

I Know What You Will Do Next Summer

Informational Privacy and the Ethics of Data Analytics

Mainz, Jakob Thrane

DOI (link to publication from Publisher):
[10.54337/aau456346604](https://doi.org/10.54337/aau456346604)

Publication date:
2021

Document Version
Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Mainz, J. T. (2021). *I Know What You Will Do Next Summer: Informational Privacy and the Ethics of Data Analytics*. Aalborg Universitetsforlag.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

I KNOW WHAT YOU'LL DO NEXT SUMMER: INFORMATIONAL PRIVACY AND THE ETHICS OF DATA ANALYTICS

**BY
JAKOB THRANE MAINZ**

DISSERTATION SUBMITTED 2021



AALBORG UNIVERSITY
DENMARK

JAKOB T. MAINZ

I KNOW WHAT YOU'LL DO NEXT SUMMER:
INFORMATIONAL PRIVACY AND THE ETHICS OF
DATA ANALYTICS

By Jakob Thrane Mainz



AALBORG UNIVERSITY
DENMARK

Dissertation Submitted

2021

Dissertation submitted: July 2021

PhD supervisor: Professor Jørn Sønderholm
Aalborg University

Assistant PhD supervisor: Associate Professor Frej Klem Thomsen
Aarhus University

PhD committee: Associate Professor Simon Laumann Jørgensen
Aalborg University (chair)

Professor Kasper Lippert Rasmussen
Aarhus University

Associate Professor Reuben Binns
University of Oxford

PhD Series: Faculty of Humanities, Aalborg University

ISSN (online): 2246-123X

ISBN (online): 978-87-7210-968-8

Published by:
Aalborg University Press
Kroghstræde 3
DK – 9220 Aalborg Ø
Phone: +45 99407140
aauf@forlag.aau.dk
forlag.aau.dk

© Copyright: Jakob Thrane Mainz

Printed in Denmark by Rosendahls, 2021

Curriculum Vitae

Prizes and Awards

2021: Winner of Res Publica's PG Paper Prize for the paper "An Indirect Argument for the Access Theory of Privacy."

Publications

1. Mainz, J. Uhrenfeldt, R. "Privacy Rights, and Why Negative Control is Not a Dead End: Reply to Munch and Lundgren." *Res Publica*. Forthcoming.
2. Mainz, J. "An Indirect Argument for the Access Theory of Privacy." *Res Publica*. Forthcoming. (Winner of the Post Graduate Paper Prize)
3. Mainz, J. Søndersholm, J. Uhrenfeldt, R. 2021. "Big Data Analytics and How to Buy an Election." *Public Affairs Quarterly*. 35(2): 119-139.
4. Mainz, J. "Are Markets in Personal Information Morally Permissible?" *Journal of Information Ethics*. Forthcoming.
5. Mainz, J. 2021. "But Anyone Can Mix Their Labor: A Reply to Cheneval." *Critical Review of Social and Political Philosophy*. 24(2): 276-285.
6. Mainz, J. Uhrenfeldt, R. 2021. "Too Much Info: Data Surveillance and Reasons to Favor the Control Account of the Right to Privacy." *Res Publica*. 27(2): 287-302.
7. Mainz, J. 2020. "Review: Lisa Herzog's 'Just Financial Markets'." *Journal of Moral Philosophy*. 17(2): 257-260.
8. Mainz, J. Søndersholm, J. "Why Some Defenders of Positive Duties Serve a Bad Theoretical Cocktail." *Journal of Global Ethics*. Forthcoming.

Education

- 2018- : PhD in Applied Philosophy, Aalborg University
- 2016-2018: MA in Applied Philosophy, Aalborg University
- 2013-2016: BA in Applied Philosophy, Aalborg University

Dansk Resumé

Hvis man ser nyheder på daglig basis, kan man næsten ikke undgå at høre regelmæssigt om, hvordan privatlivets fred trues af store tech-virksomheder og efterretningstjenesters massive indsamling og brug af personlige data om almindelige mennesker. Mange føler, at de har mindre privatliv nu, end de havde i de gode gamle 'analoge' dage, inden internettets opfindelse. Mange føler også, at den måde hvorpå deres personlige data bliver indsamlet eller brugt, 'krænker' eller forbryder sig mod deres ret til privatliv. Store virksomheder som Google og Facebook indsamler enorme mængder data om individer. Med disse data kan man ved hjælp af 'data analytics' udlede nye datapunkter om individer, som kan bruges til forudse og påvirke individers adfærd. Men er det sandt, at vi mister privatliv, når virksomheder og efterretningstjenester indsamler og udleder personlige data om os? Og er det sandt, at måderne hvorpå gør det krænker retten til privatliv?

For at kunne besvare disse spørgsmål, er det nødvendigt at vide hvad det overhovedet vil sige at have privatliv, og hvad retten til privatliv er. I denne afhandling giver jeg blandt andet (delvise) svar på disse spørgsmål. Jeg udvikler en teori om, hvordan man bedst konceptualiserer privatliv. Jeg forsvarer det synspunkt, at det at have privatliv kommer i grader, og at man har privatliv i den udstrækning at andre ikke har bestemte former for adgang til personlige information om én. Jeg udvikler

desuden en teori om, hvilke typer af handlinger, der tæller som krænkelse af retten til privatliv. Jeg forsvarer det synspunkt, at retten til privatliv krænkelse hvis, og kun hvis, man ikke har det jeg kalder 'Negativ Kontrol' over adgangen til sine personlige informationer.

Det er ofte blevet foreslået at give individer ejendomsret over personlige information om dem selv. Tanken er, at individer på den måde bedre kan beskytte sit privatliv, og at de kan købe og sælge sine personlige information på det frie marked, så det ikke kun er de store tech-virksomheder, som tjener penge på individers personlige informationer. Filosoffen Francis Cheneval har foreslået, at John Locke's teori om ejendomsret implicerer at individer rent faktisk ejer data om sig selv. Jeg forsøger at vise, hvorfor man løber ind i et dilemma, hvis man forsøger at forsvare idéen om dataejerskab vha. Locke's teori, som Cheneval gør.

Jeg forsvarer desuden det (kontroversielle) synspunkt, at det at udlede nye personlige information om folk ikke krænker deres ret til privatliv, hvis de informationer udledningen er baseret på er indsamlet på legitim vis. Hvis Facebook for eksempel på legitim vis har fået informationer om hvem dine venner er, og hvad deres politiske præferencer er, og Facebook - baseret udelukkende på disse informationer - udleder hvad *dine* politiske præferencer er, så har Facebook ikke krænket din ret til privatliv.

Sidst, men ikke mindst, påpeger jeg et demokratisk problem, der opstår når det er let at få adgang til store mængder af individers personlige data, og når man har

moderne data analytics til rådighed. Idéen er, at man med de retter mængder data, og den rette teknologi, kan købe et bestemt demokratisk valgresultat helt lovligt. Hvis man ønsker at købe et bestemt valgresultat, og man har adgang til nok personlige informationer om vælgere, så kan man med simple statistiske modeller forudse med en vis præcision, hvordan specifikke vælgere vil stemme til et givent valg. Når man først ved det, kan man indgå ansættelseskontrakter med modstanderens vælgere, der gør det umuligt for dem at stemme på valgdagen. Man kunne fx tilbyde en vælger 25 dollars mod at den pågældende vælger går rundt i den lokale park og samler skrald op på valgdagen. I et land som USA, hvor man har offentligt tilgængelige vælgerregistre, hvor man kan se *om* folk har stemt, vil man efter valget kunne tjekke, om modstanderens vælgere rent faktisk stemte. På den måde kan man med relativt få midler købe et bestemt valgresultat, helt lovligt. Dette demokratiske problem – at man kan købe et valgresultat på lovlig vis – er foranlediget delvist af at folk har begrænset privatliv, og delvist af at nye teknologier gør det muligt at udnytte dette tab af privatliv på hidtil usete måder.

English Summary

If you watch the news on a daily basis, it is difficult not to hear frequently about how your privacy is threatened by big tech-companies and intelligence services' collection and use of personal information about ordinary people like you. Many people feel that they now have less privacy than they used to have in the good old 'analog' days, before the invention of the internet. Many people also feel that the way in which their personal information is collected and used 'violates' or infringes upon their privacy rights. Big companies like Google and Facebook collect huge amounts of data about individuals. With these data, it is possible to use 'data analytics' to infer new data points about individuals, which can then be used to predict and affect the behavior of individuals. But is it true that we lose privacy when companies and intelligence services collect and infer personal data about us? And is it true that the ways in which they do so violate privacy rights?

In order to answer these questions, it is necessary to know what it even means to have privacy, and what the right to privacy is. In this thesis, I give (partial) answers to these questions – among other ones. I develop a theory of how best to conceptualize privacy. I defend the view that having privacy is a matter of degrees, and that one has privacy to the extent that others do not have certain types of access to personal information about one. I also develop a theory of which types of actions count as

violations of the right to privacy. I defend the view that the right to privacy is violated if, and only if, one does not have what I call 'Negative Control' over the access to one's personal information.

It has often been suggested to grant individuals property rights over personal information about themselves. The idea is that by granting individuals property rights over personal information, their privacy will be better protected, and they can buy and sell personal information on the free market, so that not only big tech-companies profit from individuals' personal information. The philosopher Francis Cheneval has suggested that John Locke's theory of property rights implies that individuals own personal information about themselves. I try to show how it generates a dilemma, if one attempts to defend the idea of data ownership through Locke's theory, as Cheneval does.

In addition, I defend the (controversial) view that inferring new personal information about people does not violate their privacy rights, if the information on which the inference is based are obtained legitimately. If Facebook for instance obtains legitimately information about who your friends are, and what their political preferences are, and - based on this information alone - Facebook infers *your* political preferences, then Facebook has not violated your right to privacy.

Last, but not least, I point out a democratic problem that arises when it is easy to get access to large amounts of individuals' personal data, and when one has access to modern data analytics. The idea is that with access to the right amounts of

data, and with access to the right technology, it is possible to legally buy a specific democratic election result. If one wishes to buy a specific election result, and if one has access to enough personal data about the electors, then one can with simple statistical models predict with relatively high accuracy how individual electors will vote in an upcoming election. Having this knowledge, one can offer employment contracts to the opponent's electors, which make it impossible for them to vote on Election Day. One could, for instance, offer \$25 to an elector in exchange for the elector in question to pick up trash in a local park on Election Day. In the US, there are publicly available voter registration lists, where one can see *whether* a given elector has voted. This means that one can check if a given elector actually voted, and thereby breached the contract. With a relatively low budget, one can thus buy a specific election result – completely legal. This democratic problem – that it is possible to legally buy a specific election result – is due partly to the fact that people have limited privacy, and partly to the fact that new technologies make it possible to exploit this loss of privacy in so far unprecedented ways.

Acknowledgements

Writing this section is an opportunity to reflect on the last three years, and all the help and encouragement I have received from colleagues, friends, and family. Luckily, I have been surrounded by the best team of support I could ever imagine. Before I turn to the usual list of people to thank, I wish to single out a particular person who deserves more than thanks. Jørn Sønderholm, who has served as my main supervisor throughout my PhD, has on many(!) occasions done much more to help me than what was by any means required of him. The first time I met Jørn, he was teaching formal logic at a first semester course that I was attending. I must admit that I was very intimidated by Jørn at first, and I didn't quite catch the beauty of formal logic at the time. Two semesters on, he was teaching a more advanced course in predicate logic, and I was hooked immediately. Suddenly, I could see how useful formal methods can be in philosophy, and in everyday reasoning. I initially came to Applied Philosophy to pursue an interest in applied ethics and political philosophy, but now I was contemplating becoming a logician. Since there were no more logic courses to take at Applied Philosophy, I thought I needed a different strategy, if I wanted to improve my logic skills. I'm not sure if Jørn remembers this, but one day, I knocked on his office door and asked if he had five minutes to spare. His first response was "am I your supervisor in a course or something?" I said "no", and explained that I wanted to TA his logic course if that was possible. He looked at me, very skeptically (those who

know Jørn can probably picture this if they close their eyes), and pointed to his whiteboard. “So you can do proofs like the one on the whiteboard?”, he asked. I said yes. He asked me a few basic questions about logical inference rules, and after a few minutes, he asked if I was interested in TA’ing his logic class the next semester. I accepted the offer, and I have been working closely with Jørn ever since. Along the way, I realized that becoming a logician was unrealistic without the proper formal training, which I couldn’t get at Applied Philosophy. So, I switched my main focus back to political philosophy. In the meantime, I learned that Jørn had switched his main research area to political philosophy as well many years ago. Jørn showed me the ins and outs of contemporary analytic political philosophy, and I knew that I wanted to pursue a PhD in that field if possible. Jørn jumped through countless hoops in order to help me secure funding for a PhD. One day, a few months before defending my master’s thesis, he called me and said that the private company Seluxit, who we had been in contact with for some time, had agreed to fund a part of my PhD, while the university would fund the rest (Thank you Daniel, and the rest of the crew at Seluxit!). I can’t thank Jørn enough for inspiring me to do the kind of work that I now do, for trying to teach me how to do proper analytic political philosophy, for showing me firsthand how to be a good supervisor, and for showing me how to publish well. I should add that throughout my master’s degree and most of my PhD, Jørn didn’t live in Denmark. One might think that it would be a problem doing proper supervision

over Skype for so many years. But I actually think that I have received both more and better supervision than most PhD students do. Thank you, Jørn.

I would also like to thank my secondary supervisor, Frej Klem Thomsen. Frej became my secondary supervisor when I was about a year into my PhD. Frej proved to be a good choice from the get go. I spent a week in Singapore with Frej and Jørn in the fall of 2019, where we discussed each others' papers, and gave presentations at Nanyang Technical University. The feedback I have received from Frej during the last years have been some of the most useful feedback I have ever received. Besides being an unusually sharp philosopher and supervisor, Frej is simply a genuine good-guy.

Next, I would like to thank the rest of my co-members of the Centre for Philosophy and Public Policy (C3P), Kim Angell, Rasmus Uhrenfeldt and Jens Thaysen. Kim was my office mate for two years. Kim is a Norwegian, but in spite of that, he is extremely gifted, kind, funny and friendly. It has been an absolute pleasure sharing an office with someone as brilliant and friendly as Kim. Being a Norwegian in Aalborg, it is only natural that Kim frequently wanted to explore what Jomfru Ane Gade has to offer. I have benefitted on many occasions from the fact that Kim finds the price of draught beer in Denmark ridiculously low. In the fall of 2020, Kim was offered a well-deserved position as associate professor at Tromsø University. I have learned a lot from Kim, especially that hard work and patience often pays off!

Rasmus was a fellow PhD student during the first two years of my PhD. I want to thank Rasmus for showing me how to be a good PhD student. It made my job much easier that I could ask Rasmus how to do things. I would have made so many more mistakes along the way, if it wasn't for Rasmus. As the list of papers included in the thesis shows, I have co-authored several papers with Rasmus during my PhD, which has been an absolute thrill. Rasmus has several very unique gifts. He has the gift of spotting interesting theoretical questions, and taking the arguments wherever they go. Some of our co-written papers would still be first drafts if it wasn't for Rasmus. More importantly, Rasmus has the unique gift of tracking down the best burgers in any city. He is a real foodie, but no food-snob. I have never met anyone who is as fond of McDonald's chili cheese tops as Rasmus is.

Jens was a postdoc during my PhD, and an extremely gifted one at that. If you tell Jens one sentence about a research paper you intend to write, he can quickly tell you where the argument should go, what the best objections to the argument are, and how to respond to them. And not only does Jens have the ability to see through complex arguments right away, he has also read *a lot*. If you have a question about anything from Greek mythology, over Starwars, to criminal law and political philosophy, chances are that Jens has read a book or ten about it, and remembers the relevant chapters if not the exact page number. Jens is one of the most gifted, and at

the same time down-to-earth people, I have ever met. For all practical purposes, Jens has served as a third supervisor for me.

I would also like to thank all current and former members of Center for Applied Philosophy at Aalborg University. They went from being my teachers and supervisors, to becoming my colleagues. They showed me all the different ways in which one can do proper applied philosophy. They managed to put together a new high-quality philosophy program where philosophy in all its confusing glory is represented. That in itself deserves credit. I may not have been the easiest student to convince that reading Heidegger and Derrida was worth the opportunity costs, but now I appreciate that they sometimes nudged me to get out of my comfort zone.

Thanks to my family for their endless support and encouragements. They have always supported my interests, and I am deeply thankful for that. Finally, I would like to thank my partner Sissel. She has been extremely supporting, and she has picked me up countless times, when I did not feel motivated. Sissel has made the process of writing this thesis much more enjoyable than it would otherwise have been, and she has surely helped me stay (more) sane. Thank you, Sissel!

Jakob Thrane Mainz
June 26, 2021
Aalborg

Table of Content

Curriculum Vitae	3
Dansk Resumé	4
English Summary	7
Acknowledgements	10
CHAPTER 1. Introduction.....	17
1.1. Behind the Scenes.....	17
1.2. The Common Thread of the Thesis	20
1.3. The Overall Theme of the Thesis, and Why You Should Care	28
CHAPTER 2. Methodology	38
2.1. Reflective Equilibrium.....	39
2.2. Moral Judgments	48
2.3. Thought Experiments	51
2.4. Reductio ad Absurdum	55
CHAPTER 3. What is This Thing Called Privacy?	60
3.1. The Mess of Defining Privacy	62
3.2. The Concept of Privacy and the Right to Privacy.....	78
3.3. The Wrongness of Privacy Violations	83

3.4. The Hybrid View of Privacy.....	92
3.5. Paper Summary: Too Much Info: Data Surveillance and Reasons to Favor the Control Account of the Right to Privacy	109
3.6. Paper Summary: Privacy Rights, and Why Negative Control is Not a Dead End: A Reply to Munch and Lundgren.....	110
3.7. Paper Summary: An Indirect Argument for the Access Theory of Privacy.	112
CHAPTER 4. Data Ownership	114
4.1. Paper Summary: But Anyone Can Mix Their Labor: A Reply to Cheneval	119
CHAPTER 5. Privacy & Inferences.....	122
5.1. Paper Summary: Inferences and the Right to Privacy	128
5.2. Predicting Voter Behavior	130
5.3. Paper Summary: Big Data Analytics and How to Buy an Election	135

CHAPTER 1. Introduction

1.1. Behind the Scenes

This PhD thesis is the product of three years of intense research. Looking back, the last three years have been passing by extremely fast. As anyone who has been through the process knows: writing a PhD thesis is the culmination of a steep learning curve. I would do many things differently, if I were to do it all again. However, looking at the final product, I am very happy with the result, and I am proud of the fact that I am now able to include in this thesis several papers that have been published in good journals.

When I first started working on this PhD thesis, I wanted to explore the normative implications of the advent of Internet of Things and Big Data. One of the first papers I started working on, in collaboration with Rasmus Uhrenfeldt, was a paper on the definition of the right to privacy.¹ Rasmus and I were on a research stay in Amsterdam to visit the privacy scholar Beate Rössler. Those three weeks in

¹ Throughout the thesis, when I write about a 'right to privacy' I mean a moral right to privacy as opposed to a legal right to privacy – unless specified otherwise. This is not to suggest that the arguments I make are irrelevant for legal discussions of privacy. The arguments may well have many implications for how the laws ought to protect the moral right to privacy, but these legal implications are not my primary concern in this thesis.

Amsterdam should turn out to be some of the most productive and rewarding times during my PhD. After countless rewritings, we finally got the paper **‘Too Much Info: Data Surveillance and Reasons to Favor the Control Account of the Right to Privacy’** (‘Too Much Info’ for short) published in *Res Publica*. One of the basic ideas in the paper is to borrow some of the concepts from the literature on liberty in the political philosophy literature, and apply them to the so-called Control Theory of privacy. We introduce a distinction between what we call Negative Control, Positive Control, and Republican Control, respectively. Interpreting control strictly as Negative Control solves many – if not all - of the problems that the Control Theory of privacy has faced for decades. Or at least so we argue. The idea of Negative Control should later come to play a major role in several of my papers. In **‘An Indirect Argument for the Access Theory of Privacy’** (‘An Indirect Argument’ for short), and in **‘Privacy Rights, and Why Negative Control is Not a Dead End: A Reply to Munch and Lundgren’** (‘Reply to Munch and Lundgren’ for short. Co-authored with Rasmus Uhrenfeldt), I develop the idea further, and explore the strengths and weaknesses of interpreting control as Negative Control. The idea of Negative Control picked up a bit of attention in the literature after we published *Too Much Info*. Several theorists published reply papers to us in good journals. *Reply to Munch and Lundgren* is our response to two of the critics (as the title of the paper suggests). It is one thing to be able to include in the thesis a paper that has been published in a good journal. But being able to also include a paper accepted for publication that replies to the critics

of the first paper, is very satisfying. To make things even better, *An Indirect Argument* won Res Publica's postgraduate paper prize in 2021. These are some of my biggest academic successes during my PhD, and all of them can at least partly be attributed to the idea of Negative Control that Rasmus and I developed in Amsterdam in the beginning of 2019.

Halfway through the project, I decided to change the focus of the project. I still wanted to focus on privacy, but instead of Internet of Things, I wanted to focus on issues related to data analytics more generally. This quickly spawned a new paper with Rasmus Uhrenfeldt. Rasmus was working on a project on the secret ballot, and we decided to write a paper on how data analytics threatens the secrecy of the vote. Later, Rasmus, Jørn and I decided to write a spin-off paper titled '**Big Data Analytics and How to Buy an Election**' ('How to buy an election' for short), where we demonstrate how it is possible to legally buy an election in the US. The paper is now published in Public Affairs Quarterly. While Rasmus, Jørn and I were working on the paper, I was also pursuing several other projects. One of them was the paper '**Inferences and the Right to Privacy**' ('Inferences' for short). In this paper, I wanted to answer the broad question of whether the use of data analytics to infer personal information about individuals violates their privacy rights. As we shall see later, the answer I give to this question may conflict with what I argued in *Too Much Info* and in *Reply to Munch and Lundgren*.

In the beginning of 2020, I returned to Amsterdam for a longer research stay. Unfortunately, I had to rush home to Denmark way too early due to the outbreak of the Covid-pandemic. Those six weeks were extremely productive, though. I wrote the paper **‘But Anyone Can Mix Their Labor: A Reply to Cheneval’** (‘Reply to Cheneval’ for short) in the first few weeks while I was in Amsterdam, and the paper was accepted for publication before I returned home to Denmark.

In addition to the six papers included in the thesis, I have been working on a (probably too) large number of papers pretty much unrelated to the topic of this thesis. I have included them all on the list in the next section. It has been a pleasure working on these side projects, and several of them have been published in good journals. At the beginning of my PhD, I was worried that I was not able to write enough publishable papers to have a thesis to submit. Now I can see that this worry was unfounded, since I have written or co-written a total of 15 papers during the last three years. In hindsight, I should probably have focused more on quality and less on quantity. But what is the PhD process if not an opportunity to learn?

1.2. The Common Thread of the Thesis

This thesis consists of five introductory chapters and six research papers. The chapters introduce a range of themes related to the research papers. Each thematic chapter

contains a summary of one or more research papers related to the theme of the chapter in question. The six research papers follow each other consecutively at the end of the thesis. Here is a list of the research papers included in this thesis:

1. Mainz, Jakob. Rasmus Uhrenfeldt. 2021. Too Much Info: Data Surveillance and Reasons to Favor the Control Account of the Right to Privacy. *Res Publica*. 27(2): 287-302.
2. Mainz, Jakob. Rasmus Uhrenfeldt. Privacy Rights, and Why Negative Control is Not a Dead End: A Reply to Munch and Lundgren. *Res Publica*. Forthcoming.
3. Mainz, Jakob. 2021. An Indirect Argument for the Access Theory of Privacy. *Res Publica*. Forthcoming. (Winner of Res Publica's PG Paper Prize, 2021).
4. Mainz, Jakob. 2021. But Anyone Can Mix Their Labor: A Reply to Cheneval. *Critical Review of International Social and Political Philosophy*. 24(2): 276-285.
5. Mainz, Jakob. Inferences and the Right to Privacy. *Journal of Political Philosophy*. Under Review.
6. Mainz, Jakob. Rasmus Uhrenfeldt. Jørn Sønderholm. 2021. Big Data Analytics and How to Buy an Election. *Public Affairs Quarterly*. 35(2): 119-139.

Here is a list of the research papers written during my PhD that are *not* included in this thesis:

7. Mainz, Jakob. 2021. Are Markets in Personal Information Morally Permissible? *Journal of Information Ethics*. Forthcoming.
8. Sønderholm, Jørn. Jakob Mainz. Why Some Defenders of Positive Duties Serve a Bad Theoretical Cocktail. *Journal of Global Ethics*. Forthcoming.
9. Munch, Lauritz. Jakob Mainz. To Believe, or not to Believe – That is Not the (Only) Question: A Hybrid View of Privacy. *Synthese*. Under Review.
10. Mainz, Jakob. Who Cares If We Can Trust Medical AI? *Journal of Medical Ethics*. Under Review.
11. Mainz, Jakob. Normative and Descriptive Theories of Privacy: How to Solve the Parent/Macnish Dilemma. *Journal of Social Philosophy*. Revise & Resubmit.
12. Sønderholm, Jørn. Jakob Mainz. Driving it Home: Why Busing Electors to the Polling Station on Election Day is an Instance of Paying People to Vote. *Law and Philosophy*. Under Review.

13. Mainz, Jakob. 2020. Review: Just Financial Markets? Finance in a Just Society – Written by Lisa Herzog. *Journal of Moral Philosophy*. 17(2): 257-260.
14. Mainz, Jakob. Rasmus Uhrenfeldt. Not So Secret After All: How Big Data Threatens the Secret Ballot and What (not) to do About It. Draft.
15. Paaske, David. Jakob Mainz. If You Polluted, You're Included: The All Affected Principle and the Democratic Enfranchisement of Polluters. Draft.

In this section, I will briefly explain how the papers included in this thesis fit together and - perhaps more importantly - how they do not fit together. I consider several of the papers *not* included in this thesis to be important for the overall story I want to tell. However, according the formal rules, the thesis may not contain more than six papers. Thus, I will focus only on the common thread of the papers that actually ended up being included in the thesis, namely the papers 1-6 on the list above. Throughout the thesis, I will however make occasional references to some of the papers that are *not* included in the thesis.

As the title of the thesis suggests, the overall theme of the thesis is that of informational privacy, and the ethics of data analytics. This theme is very wide in some ways, and very narrow in others. For instance, the part about data analytics is very wide, in the sense that the concept of data analytics itself is very wide. Many

different types of fields and practices (such as supervised/unsupervised machine learning, deep neural networks etc.) are conventionally categorized under the umbrella-term of ‘Artificial Intelligence’. The term ‘data analytics’ is purposely underspecified, because nothing in this thesis forces me to specify it much further. However, I take (at least the modern versions of) ‘data analytics’ also to be categorized under the umbrella-term of Artificial Intelligence. The key features of data analytics - broadly understood - that are relevant for this thesis is the ability to accurately correlate pieces of information and the ability to accurately infer ‘new’ pieces of information – especially personal information - which can be used to predict the behavior of individuals. When I talk of ‘data analytics’ or ‘big data analytics’, I refer to types of data analytics that have at least these basic features.

The part about informational privacy is in a way very narrow, given that it leaves out many relevant and interesting aspects of privacy. Part of the reason for focusing on informational privacy is that this aspect of privacy is presumably the most relevant one in the context of data analytics. But, this limited scope is also motivated by the consideration that what I take to be the most thorough and analytically rigorous part of the privacy literature is exactly the part of the literature that focuses on informational privacy. Most of the thesis is concerned with the *right* to privacy (or at least normative aspects of privacy), rather than the *concept* of privacy as such. However, in order to properly understand what the right to privacy is, and how it

works, it is useful if not necessary to get a proper grasp of what the concept of privacy is. Thus, some of the papers – and some of the chapters in the thesis - also have a partial focus on the *concept* of privacy.

With these brief preliminary remarks out of the way, let us now turn to the common thread of the papers included in the thesis:

In *Too Much Info*, Rasmus Uhrenfeldt and I argue that there is at least a pro tanto reason to favor the so-called Control Theory of the right to privacy over the rival Access Theory. We bracket the question of how to conceptualize privacy, and focus on what the right to privacy is, and in particular, which types of actions count as violations of the right to privacy. As mentioned earlier, after publishing *Too Much Info*, several theorists published responses to the paper. Two of the responses were published in Res Publica.

As the title of the paper suggests, *Reply to Munch and Lundgren* is Rasmus Uhrenfeldt and my response to the objections raised in two of the replies to *Too Much Info*. In *Reply to Munch and Lundgren*, the issue of the concept of privacy plays a crucial role, and we suggest a new account of how the concept of privacy relates to the right to privacy. We also try to improve the definition of Negative Control that we first introduced in *Too Much Info*.

In *An Indirect Argument*, I pick up the idea of Negative Control again, and argue that while this idea is promising in the context of the right to privacy, it is not promising in the context of the concept of privacy, because it collapses the Control Theory into the rival Access Theory. The notion of Negative Control allows the control theorist to avert all the classic objections against the Control Theory. But this is not a win for the control theorist, because the non-normative version of Negative Control sneaks in notions of access in a way that makes it coextensive with the Access Theory. The three papers mentioned above constitute the part of this thesis concerned with the theoretical issues related to privacy, and especially the right to privacy.

Two of the remaining papers, *Inferences* and *How to Buy an Election* are more concerned with some of the implications of the arguments defended in the theoretical papers on privacy described above. In *Inferences*, I describe how data analytics is used to infer ‘new’ personal information about individuals. I defend the view that such inferences do not violate the privacy rights of individuals, if the information that the inferences are based on are themselves obtained legitimately. In *Inferences*, I bracket the dispute between the Control Theory and the Access Theory, and focus on the question of whether inferring information based on information that is obtained legitimately count as a violation of the right to privacy. I do not discuss whether the conclusion defended in *Inferences* is compatible with the idea of Negative Control defended in *Too Much Info*, and in *Reply to Munch and Lundgren*. My main

worry is, however, that interpreting control as Negative Control implies that the view I defend in *Inferences* is incorrect, and vice versa, or at least that there are some problematic tensions between those two views.

In *Reply to Cheneval*, I turn to the idea of protecting individuals' privacy rights by granting them property rights over personal information that pertains to them. The idea of 'data ownership' has been a hot debate in economics, law, and in the general public for a long time. But scant attention has been paid to this idea in the philosophy literature. A recent exception to this is a paper by Francis Cheneval published in *Critical Review of International Social and Political Philosophy*. Cheneval argues that Locke's theory of property rights imply that people own data that pertains to them. In *Reply to Cheneval*, I try to explain why Cheneval's argument runs into a dilemma.

Finally, in *How to Buy an Election*, Jørn Sønderholm, Rasmus Uhrenfeldt, and I discuss a democratic problem that arises when individual voters lose privacy over information that make it possible to use data analytics to predict how they are going to vote in an upcoming election. Our main finding is that in an electoral system that has publicly available voter registration lists (like the one in the US or the UK), individual voters' information can be exploited in a way that makes it possible to legally buy an election. This paper does not discuss the right to privacy. Rather, it

discusses other normative problems that arise when many individuals lose privacy in a non-normative sense.

This concludes the brief summary of how the papers relate to each other. I elaborate on these relations throughout the thesis. In the next section, I will briefly introduce the overall theme of the thesis, and explain why I think it is important.

1.3. The Overall Theme of the Thesis, and Why You Should Care

Most people probably have some intuitive ideas of what privacy is, and why it is important. But most people probably also have various *conflicting* ideas of what privacy is and why it is important. This is no surprise. As will become evident throughout the thesis, the concept of privacy covers many different things, and these things are not always obviously related. Worse, when they are related, they are not always compatible. This is part of the reason why a significant part of the thesis is dedicated to defining and clarifying what privacy is, and especially what the right to privacy is. As we shall see, this is no easy task.

Although it is notoriously difficult to define privacy, common sense conceptions of the term take us a long way. Despite having many different – and perhaps conflicting – ideas of what privacy is many people nevertheless seem to worry

about their privacy. If you turn on the news on a daily basis, you will often see stories about how companies or governments surveil people, how big tech companies collect huge amounts of data and predict the behavior of people. Every time we go online, use a credit card, drive a car, watch TV, use a fitness tracker, walk around in the city, or cross a national border, our behavior is tracked minutely one way or another. We all leave digital trails wherever we go.

Scandals like the one involving Cambridge Analytica and Facebook, or the one involving Edward Snowden's revelations of how the NSA surveils innocent people around the world, have made many ordinary non-tech-savvy people aware that both private companies and governments collect huge amounts of data about all of us, and that they often use the data in objectionable ways. Many people know this by now – even if they do not fully grasp the extent to which this is so - or even if they do not know the technical details of how it works.

When asked, people often report that they are very concerned about their privacy (Madden & Rainie, 2015). Nevertheless, many people do not act accordingly (Hargittai, E., Marwick, A., 2016). On the one hand, they worry about what their data is used for, and they do not like the idea of companies or states having access to all sorts of information about them. But talk is cheap. When we study how people actually behave, it turns out that most people do next to nothing to protect their privacy. This phenomenon is called the 'Privacy Paradox' (Barnes, 2006), and it is a well-studied

and well-described phenomenon that occurs consistently across national borders, age groups, cultures, etc.

There are many competing explanations for the Privacy Paradox. One explanation is that people do not sufficiently understand the risks associated with sharing their data. Another explanation is that people are not sufficiently informed about which privacy-protecting measures they can take (Hargittai, Eszter & Litt, 2013). A third explanation is that despite being aware of all the relevant risks and all the relevant counter-measures, the social advantages of disclosing information on social media etc. are simply worth the loss of privacy (Taddicken, 2014). A fourth explanation is that many people believe that despite the existence of certain privacy-enhancing measures that one can take, there is not really much one can do to effectively avoid the collection of personal data (Hargittai, E., Marwick, A., 2016).²

There is much to be said about each of the explanations. In order to motivate the overall theme of the thesis, I will briefly say something to the effect that those who worry about their privacy may in fact be more justified in doing so than most of them

² See (Barth & Jong, Menno D. T. de, 2017) for a systematic literature review on the Privacy Paradox.

probably realize. I do not consider myself a ‘privacy alarmist’, but neither do I consider myself naïve. There are indeed real and hard-to-deny reasons to be worried about one’s privacy – and in fact also about the privacy of others.

As I shall discuss later, some theorists think that the right to privacy can be violated even when the loss of privacy in question – such as a hacker’s access to your health information – never materializes in any subsequent harm. For instance, if the hacker gains access to your health information and discovers that you have cancer, then the hacker has violated your right to privacy. This is so, even if the hacker never shares the information with others, never uses the information to blackmail you, or anything of that sort. Similarly, it is a violation of your right to privacy if your neighbor sets up a camera over the hedge and spies on you having sex. This is so, even if no one else but your neighbor watches the sex tape, even if the neighbor never tells anyone about the sex tape, or anything of that sort. Even if you never find out that your neighbor recorded you, your right to privacy is still violated. Or at least so many privacy theorists (and common sense morality) suggest.

Nevertheless, it is not uncommon to hear people say things like “what’s the harm of a privacy violation?”, or “people can look all they want at my data, as long as they don’t misuse it”. For what it is worth, anecdotal evidence tells me that people who have such sentiments are often difficult to convince otherwise. Perhaps rightly

so. But, it is probably not necessary to convince them otherwise, because many privacy diminishments in fact *do* have harmful consequences.

Many of these harmful consequences are not even faced by the data subjects themselves, but by others. The data that we share about ourselves, when we create an account on social media, when we use our credit card in the local grocery store, or when we pay back a loan in the bank, is data that affects the lives of others. Even if you do not care about your own privacy, and happily share all sorts of personal information with others, you should know that not caring about your own privacy makes it difficult for others to care about theirs. This is not to suggest that it is not ultimately up to you to decide what information you share about yourself. After all, there are many actions that potentially harm others that are nevertheless permissible to do (think for example of the action of leaving a romantic partner who you do not love anymore, but who still loves you). But, the fact that sharing your data potentially has harmful consequences for others, at least gives us a *pro tanto* reason to think that it may sometimes be wrong to share certain pieces of information about yourself. And, it gives us a *pro tanto* reason to think that the subject of privacy, and perhaps especially informational privacy, is a morally important one.

The consequences that others may face when you decide to share your data, can be of many different types. Sometimes the consequence is simply that other people's privacy is affected as well. If you share some information about yourself on

Facebook, and Facebook knows that you are friends with Alice and Bob, then Facebook may be able to couple this information with other pieces of information, and infer ‘new’ information about Alice and Bob. This may be information that Alice and Bob never wanted Facebook to have. In effect, the level of privacy that Alice and Bob enjoy is sometimes partly a function of the level of privacy that you enjoy. Solon Barocas and Karen Levy have recently called this phenomenon ‘privacy dependency’ (Barocas & Levy, 2020). Alice and Bob’s privacy depends in part on your privacy.

Barocas and Levy describe three different types of privacy dependencies: Similarity-based dependency, difference-based dependency, and tie-based dependency. In a similarity-based dependency, the inferences are based partly on the fact that you share certain similarities with Alice and Bob. In a difference-based dependency, the inferences are based partly on the fact that you do *not* share certain similarities with Alice and Bob. And finally, in a tie-based dependency, the inferences are based partly on the fact that you have certain social ties to Alice and Bob. It could

be, for instance, that they are your siblings, your colleagues, or something similar (Barocas & Levy, 2020).³

A famous example involving privacy dependencies is the following: The supermarket company Target allegedly predicted that a sixteen year old girl was pregnant before her dad knew about it (Hill, 2012). Target collected data about what products pregnant customers tend to buy, and then trained a machine learning model to predict if specific customers were pregnant, in order to send them advertisement for pregnancy-related items. One day, an angry man walked into the local store and complained that his sixteen-year-old daughter was receiving pregnancy-related ads at their home address. Later, the man learned that his daughter actually was pregnant, and he apologized to Target.

I shall return to the idea of privacy dependencies in the paper *Inferences* (although I do not use this terminology in that paper). In *Inferences*, I defend the view that inferences of other people’s personal information do not in themselves violate privacy rights, if the information on which the inferences are based are obtained

³ For discussions of concepts related to that of ‘privacy dependencies’, see e.g. (Fairfield & Engel, 2015).

legitimately. For instance, if the data on which Target trained their model was collected legitimately, then my argument implies – perhaps controversially - that Target did not violate the privacy rights of the sixteen-year-old girl. I discuss and reject the objection that this is not so, if the inference in question is other-regarding in the sense involved in privacy-dependencies.

Now, privacy dependencies are surely interesting and important. But the consequences of privacy losses or privacy diminishments can also have far more drastic consequences for others. For instance, the data that you share can have consequences for whether or not other people are subjected to online voter manipulation (Susser, Roessler, & Nissenbaum, 2019), whether or not they can get a loan in the bank (Turkson, Baagyere, & Wenya, Sep. 2016), get the job they applied for (Raghavan, Barocas, Kleinberg, & Levy, 2020), get released from prison on parole (Castro, 2020), get accepted at a good college (Kuyoro, Goga, Awodele, & Okolie, 2013), get an affordable premium on their insurance (Noorhannah & Jayabalan, 2018), or even whether or not they are likely to have encounters with the police (Meijer & Wessels, 2019).

Private companies and governments alike use the data that you share to train predictive machine learning models in order to make all sorts of important decisions about you and others. If the insurance company discovers that ‘people like you’ tend to suffer from all sorts of lifestyle-related diseases early in life, then your health

insurance premium may go up. If the bank discovers that people like you tend to default on their loans, then your loan application may be rejected. If the company discovers that people like you tend to perform badly in certain jobs, then your job application may be rejected. If the courts discover that people like your tend to recidivate, then your parole application may be rejected. If political parties discover that people like you tend to vote for the opponent party, then they may try to use clever manipulation tactics to persuade you to vote differently. And so on. There are countless examples of how data about individuals is being used on a grand scale to make inferences about other people, which again is used to predict their behavior, and make important decisions based on these predictions. Of course, none of this is new. Banks have for a long time been trying to figure out if people like you are likely to default on a loan. Employers have for a long time been trying to figure out if people like you are likely to be a good employee. And so on. What is new is the scale, speed, and accuracy at which all this can be done due to the advent of machine learning. Moreover, the affects all of this has on your privacy is completely unprecedented. The amounts of personal data that are collected and used in order to make accurate predictions are far vaster than they used to be before the advent of machine learning.

I am not suggesting that you should only care about your privacy, because your personal information is used to make important decisions that affect the lives of others. Privacy is important for many different reasons. But even if you do not care

about privacy for *its own sake*, perhaps you should care about some of the *consequences* of losing privacy. Or, perhaps you should care about your privacy because companies at the moment monetize on your data, while you get no cut yourself. Whether or not you care about privacy at all, and regardless of your reasons for caring about privacy if you do, I hope that this thesis will at least help you get a better grasp of what privacy is all about, and what some of the implications of lacking privacy can be. Let us now turn to the methodology used in this thesis.

CHAPTER 2. Methodology

In this methodology chapter, I will do two things. First, I will explain *how* particular components of what we might call the ‘standard method’ of contemporary analytic political theory is normally used. Second, for each methodological component, I will show how I have made use of it in this thesis.

Before turning to the description of each of the methodological components, I want to highlight what type of research questions I am addressing in this thesis. The thesis deals primarily with normative questions like how we *ought* to define privacy rights, how we *ought* to block the possibility of using data analytics to buy an election, and so on. These are normative questions. Nevertheless, the aim of the thesis is to get to the truth of the matter. To many, this will sound somewhat controversial. How can normative questions about what we *ought* to do be the subject of scientific inquiry? My aim of this methodology chapter is not to refute skepticism about treating normative questions in a scientific way. Rather, I will try to explain *how* the method works in general, and how I have used it in particular.

Throughout the thesis, I assume that the metaethical position of Moral Cognitivism is true. That is, normative propositions are truth-apt; they can be either true or false. Without the assumption of Moral Cognitivism, the method I am about to describe does not work the way it is usually said to work. If normative propositions

are not truth-apt, then no argument containing normative premises can be valid. And, if no such argument can be valid, it cannot be sound either (because the definition of a sound argument is a valid argument with true premises). Similarly, if Moral Cognitivism is false, then the standard method of ‘Reflective Equilibrium’ (RE) does not work either. With the assumption of Moral Cognitivism in the back of our minds, let us now turn to the method of RE.

2.1. Reflective Equilibrium

Many political theorists seem to think that the method of RE is the standard method of contemporary analytic political theory (Knight, 2017: 46); (Sinnott-Armstrong, Young, & Cushman, 2010: 246). John Rawls famously coined this method (Rawls, 1971), but it was commonly used long before Rawls.

The idea of RE is relatively simple. We all have a set of moral judgements about particular cases. For example, most of us can agree that under normal circumstances, Smith is doing something wrong if we find him torturing a baby for fun, peeping into the ladies’ room, or stealing a car. The overall goal of RE is to balance these judgments with the set of moral principles that we also believe to be true. For example, we might believe that it is always bad to unjustifiably cause harm to other people. This principle is easy to square with the judgment about Smith’s behavior. If the principle is true, then it is clear that Smith ought not to torture this

particular baby for fun, let alone any other baby. If the principle under consideration had not been sitting well with the moral judgment in the particular case, then we would need to go back and forth, revising either our principle or our judgment, until there is no conflict between the two. If we can reach a condition where all of the principles we believe to be true, and all our judgments about particular cases, fit together in a coherent way, then we have reached a RE. It is therefore somewhat misleading to talk about RE as a method in itself. It is probably more accurate to talk of RE as the end-condition that obtains when we have balanced our particular judgments with our principles.

The type of RE that emerges when the principles and the particular judgments fit together in a coherent way is often called ‘Narrow RE’. In addition to the Narrow RE, we can also try to reach a ‘Wide RE’. Narrow RE is reached when (a) a set of considered moral judgments⁴, and (b) a set of moral principles, fit together in

⁴ To say that a moral judgment is ‘considered’ can mean at least two things (Rawls, 1971: 47). First, it can mean that the judgment is made without undue influence, which

a coherent manner. Wide RE is reached, when (c) a set of relevant background theories is added to the system, and coherence between (a), (b), and (c) still upholds. If the background theories do not immediately fit together with the moral judgments and the moral principles, then the three of them need to be reconsidered and revised again, until coherence emerges (Daniels, 1979: 258). What we mean by background theories here can be understood very loosely, as any theory that is in any way relevant for the truth of (a) and (b). We might, for example, consider utilitarianism as a background theory. We might, *prima facie*, think that - as an instance of (a) - it is wrong for Smith to torture the baby for fun. We might also think - as an instance of (b) - that the principle which holds that one ought not to unjustifiably harm others is true. Further, we might think - as an instance of (c) - that utilitarianism is the correct moral theory. At first, it may seem as if these three propositions cannot be coherent. If that is the case, then we need to revise at least one of them. But if we look closer, we see that the three propositions are coherent. For example, it might be that our particular

may make the judgment unreliable. Examples of this, could be making obviously biased judgments of the sort famously discussed by Daniel Kahneman (Kahneman, 2013). It could also be judgments that are vulnerable to so-called ‘debunking arguments’, such as judgments that are unreliable for evolutionary reasons (Singer, 2005). Second, it can mean that the person making the judgment is relatively confident that it is correct (Rawls, 1971: 47). See (Knight, 2017: 47) for a good critique of this ‘confidence constraint’.

judgment about Smith torturing the baby is compatible with, and explained by, the principle that one ought not unjustifiably harm others. This principle might then be explained by - or at least be compatible with - the background theory of utilitarianism. I do not claim that this is in fact the correct Wide RE to reach here. I think that if we added certain other background theories, or considered certain other moral judgments, we would come to reject utilitarianism. However, the point is merely to demonstrate the technical process of reaching Wide RE.

This point is crucial: The background theories are not written in stone, just like the moral judgments in (a), and the moral principles in (b), are not written in stone. We do not add the background theories merely to check if the Narrow RE we reached was correct. Rather, we might need to adjust and possibly reject the background theories as well. Norman Daniels puts the adequate role of background theories in Wide RE like this:

The background theories in (c) should show that the moral principles in (b) are more acceptable than alternative principles on grounds to some degree independent of (b)'s match with relevant considered moral judgments in (a). If they are not in this way independently supported, there seems to be no gain over the support the principles would have had in a corresponding narrow equilibrium, where there never was any appeal to (c). Another way to raise this point is to ask how we can be sure that the moral principles that

systematize the considered moral judgments are not just “accidental generalizations” of the “moral facts,” analogous to accidental generalizations which we want to distinguish from real scientific laws. (Daniels, 1979: 259).

The background theories are thus necessary to render probable that the principles do not just happen by coincidence to track the particular moral judgments. If the moral judgments in (a), the principles in (b), and the background theories in (c) all form one coherent system, then at least it becomes less probable that the principles are just accidental generalizations of the moral judgments.

I have now painted - in very broad strokes - the common method of RE in analytic political theory. Some readers may think that my description is underspecified, or that I have put too much, or too little weight on certain components of the method. I believe, though, that what I have described is sufficient to understand what is going on in the papers that make up this thesis.

Let us now look at an example of how I make use of the method in this thesis. In *Too Much Info*, Rasmus Uhrenfeldt and I follow the method of RE stringently. Not only do we use the method to arrive - behind the scenes - at the conclusion we want to defend. We also describe each step of the process explicitly in the paper.

Given our judgments about particular cases, our point of departure in the paper is that the following version of the Control Theory of the right to privacy is at

least prima facie plausible, but that it is not able to explain certain moral judgments that we find obviously true:⁵

CA1: For any agent A to have her right to privacy violated, there is a necessary and sufficient condition that must be satisfied: Agent A has involuntarily lost control over unwanted access to personal information P about agent A. (Mainz & Uhrenfeldt, 2021: 289).

The first point we make in the paper is that CA1 cannot accommodate the judgment that no violation of the right to privacy seems to occur in the following case:

Too Much Info #1: Suppose that Smith and Jones are co-workers. Smith likes to share personal information about his sex life. One day, as Smith is about to tell Jones something personal again, Jones simply puts his fingers in his ears before Smith starts talking. Smith finishes his story anyway. (Mainz & Uhrenfeldt, 2021: 294).

⁵ In the paper, we call it the ‘Control Account’ (CA) rather than the ‘Control Theory’.

In Too Much Info #1, it seems undeniably true that Jones does not violate Smith's right to privacy by putting his fingers into his ears when Smith is about to tell him something private. Nonetheless, at least on one interpretation of the word 'control' common in the literature, Jones does in fact violate Smith's right to privacy in this case. To spell out what interpretation we are talking about, we present three types of control:

Positive Control: Agent A enjoys Positive Control over the access to relevant information P, if, and only if, A tries (or could try) to give agent B actual access to P, and succeeds.

Negative Control: Agent A enjoys Negative Control over access to relevant information P, if, and only if, A is capable of preventing agent B, who attempts to access, from accessing P.

Republican Control: Agent A enjoys Republican Control if, and only if, agent B does not have the ability to get access to relevant information P about A. (Mainz & Uhrenfeldt, 2021: 293).

The point is that if control is interpreted as Positive Control, then it has the counterintuitive implication that Jones violates Smith's right to privacy when he puts his fingers in his ears, since then Smith no longer has control - in the positive sense - over his personal information.

In order to avoid this implication, we suggested - following RE - that the Control Theory should be reformulated in the following way:

CA2: For any agent A to have her right to privacy violated, there is a necessary and sufficient condition that must be satisfied: Agent A has involuntarily lost *negative* control over unwanted access to personal information P about agent A. (Mainz & Uhrenfeldt, 2021: 292).

We then move back and forth between different formulations of the Control Theory, and particular judgments about particular cases, until we reach CA4:

CA4: For any agent A to have her right to privacy violated, there is a necessary and sufficient condition that must be satisfied: Agent A has involuntarily lost Negative Control over the access to personal information P about agent A, due to action(s) of agent B, of which B is responsible. (Mainz & Uhrenfeldt, 2021: 298).

We do not claim in the paper that Narrow RE obtains when we reach CA4. We might be able to keep reformulating specific parts of the CA4 if we considered more cases. In *Too Much Info*, however, the point is not to show that CA4 is the ultimate principle of privacy rights, but rather that it reaches a more stable equilibrium than the competing Access Theory does.

Finally, in the paper we consider a set of background theories, in order to achieve Wide RE rather than just Narrow RE. In particular, we consider a set of theories about rights in general, and conclude that regardless of which of these background theories we consider, it does not force us to reconsider the principle CA4, nor the moral judgments about the particular cases:

When we say that a person has a right to privacy, we do not subscribe to any particular theory of what it means to have a right to something. All our arguments are compatible with all of the most common theories of rights. For example, according to the interest theory of rights, the function of a person's right to privacy is that having such a right furthers her interests. According to the will theory of rights, on the other hand, the function of a person's right to privacy is to give that person control over the duties of other persons with regards to her privacy. Since nothing in our arguments hangs on which account of rights is the correct one, we will remain agnostic about this. However, we will assume—uncontroversially—that a right to privacy is a waivable, non-absolute right. (Mainz & Uhrenfeldt, 2021: 288).

In this section, I have briefly explained how the method of RE works, and I have provided examples of how I use this method in the paper *Too Much Info*. In the next section, I will explain what we mean by the term 'moral judgments', and what role such judgments play in the method.

2.2. Moral Judgments

One of the most controversial components of the method of RE is the use of moral judgments - or 'moral intuitions' as they are sometimes called - about particular cases. Perhaps the most controversial thing about the use of moral judgments is the epistemic weight we often ascribe to them (Huemer, 2005). One might accept that normative propositions have truth-values. But one might still insist that even if that is the case, we have no (or very limited) epistemic access to normative truths (McMahan, 2013). In the positive sciences, we use our senses to make observations about the physical world. But we have no special moral sense which allows us to make 'moral observations'. We cannot perceive morality the way we perceive physical objects. The methods used in the positive sciences are in a way 'foundationalist', in the sense that they treat observations as being more or less foundational: The observations are not justified only if they can be explained by theory, the theory is judged by its ability to explain the observations. In the positive sciences, we rarely distrust the observations when they do not fit our theory (at least not in idealized positive sciences!). The justificatory relation between theory and observation in the positive sciences is not analogous to the justificatory relation between theory and moral judgment in moral theory. In moral theory, we move back and forth until we reach a wide RE. In the process, we relatively often revise our moral judgments. In the positive sciences, however, there is less moving back and forth, because observations are treated as more

or less foundational. Nevertheless, the method of RE assumes that our moral judgments about particular cases carry at least *some* justificatory weight. Otherwise, the method of RE would be completely redundant, if not impossible to carry out. It is important to note, however, that this characteristic of the method of RE does not imply that the method of RE is entirely a ‘coherentist’ theory, even though many theorists have claimed that it is. As Carl Knight has recently pointed out, the method of RE contains both coherentist- and foundationalist elements (Knight, 2017: 50). We can treat certain moral judgments or specific moral principles as more or less foundational, and still move back and forth between non-foundational judgments, principles, and background theories, until we reach a coherent set of beliefs.

It seems that the method of RE is dependent at least to some extent on the reliability of our moral judgments. My own view – and this has no particular bearing on how I use the method of RE in my papers – is some version of Intuitionism. By that, I mean that moral judgments carry *pro tanto* justificatory weight. Absent good reasons to believe otherwise, I should trust that my considered moral judgment in a particular case is correct (Huemer, 2005). Moral judgments are much less mystical than what they often get credit for. It is true that the objects of moral judgments are different from the objects of non-moral judgments. It is also true that we can use our senses to make observations about, say, the colors of the rainbow, but that we cannot use our senses to make observation about, say, the morality of torturing a baby for

fun. However, it does not follow from this that the non-moral judgments themselves are of a different type than the moral judgments. The objects of the judgments are different, but it is not obvious that the judgments themselves are qualitatively different. Much of the skepticism about moral judgments may stem from the ontological difference in the objects being judged, and not from some epistemic difference between two distinct types of judgments. That said, there is certainly room for skepticism about our judgments in general, and perhaps especially about moral judgments in particular. As many theorists have pointed out, several factors can render moral judgments more or less reliable, such as evolutionary dispositions, general cognitive biases, nutrition, physical energy levels, mental stability, and so on (Singer, 2005). Note, however, that the same is often true of non-moral judgments. You can certainly also be evolutionarily disposed or cognitively biased when making certain non-moral judgments. Perhaps the primary difference, which gives us reasons to be more skeptical about moral judgments than non-moral judgments, is that we are more frequently biased when making moral judgments than we are when making non-moral judgments.

Luckily, just like a PhD thesis in physics does not need to establish *why* we can trust our sensory perceptions, I need not establish *why* we can trust moral judgments. The important thing is that *if* we accept that moral judgments have *some* justificatory weight, then the method of RE takes us quite far in our normative inquiry,

at least in combination with an assumption of the truth of Moral Cognitivism, and a general demand for logical consistency.

2.3. Thought Experiments

One need not read much analytic political theory before one stumbles upon a thought experiment, often a very outlandish one (Bunzl, 1996). But what is the role of these thought experiments in analytic political theory? Generally speaking, thought experiments are used to ‘pump’ our judgments about some issue (Brendel, 2004). These judgments can be normative - like the ones discussed in the previous sections - but they can also be non-normative. For example, we sometimes use thought experiments to make conceptual points. We might ask what the scope of the concept of ‘privacy’ is. Then we might run through a series of thought experiments in order to see if privacy obtains in those specific cases. Charles Fried famously asks the reader to imagine a person stranded on a deserted island, where no one can ever see him or hear him. Fried then asks the reader if that person enjoys privacy (in other words, if that person is in a condition of privacy). To Fried, the obvious answer is “no”. As long as the person on the island is not able to grant or deny anyone access to his personal matters, it is ironic to speak of privacy (Fried, 1968: 482). Here, it is important to note that Fried is not discussing the right to privacy, but rather the condition of privacy. The thought experiment is therefore not meant to pump any moral judgment, but rather

to make a conceptual point. I do not share Fried's judgment. I think it is perfectly sensible to say that the person on the island enjoys privacy. Differences in judgments - like the difference between Fried's conceptual judgment and my conceptual judgment about the man on the island - may partly explain why we subscribe to different moral principles, and perhaps in the end reach different Wide RE.

Very often, though, we do use thought experiments to pump *moral* judgments. Thus, the thought experiments fit into the method of RE since thought experiments are very useful ways of exploring what our judgments are, and possibly how strong particular judgments are. There can be several reasons why we might want to construct a thought experiment instead of referring to actual cases that have actually occurred in the real world. First, the scenario we need may not, as a contingent fact, have occurred in the real world yet, or we might not be able to know or document that it has occurred. Second, sometimes cases from the real world come with a lot of 'noise'. That is, it is often easier to extract the morally relevant features, and only the morally relevant features, in a thought experiment, than it is in a real world case.

Some theorists claim that it is better to rely on real world cases when such are available, or that there is a limit to how outlandish the thought experiments may be, even if no real world cases are available to us. Jakob Elster has famously argued that we cannot trust our moral judgments about very outlandish cases. He argued against 'Conceivabilism', and in favor of 'Realism':

Conceivabilism: As long as a case is conceivable, it is legitimate to use it to elicit intuitions for testing moral principles

Realism: Only cases which could plausibly occur given the world as it is should be used to elicit intuitions. (Elster, 2011: 242).

Elster's point is not that we can only rely on our moral judgments about cases from this world, and not on our judgments about cases from other possible worlds. Rather, his point is that if the case is very outlandish, then we cannot reliably trust our judgment about these cases, since we simply lack the capacity to apply our intuitive moral competences to such cases. Elster's main interlocutor is Kasper Lippert-Rasmussen, who is notorious for making use of very outlandish thought experiments like the following:

'Two hundred legs and arms'. Suppose, for instance, that people are born with huge bodies they can barely move, bodies with two hundred legs and arms. At any given moment, they can at best sense and control 1 percent of their bodies, although they can readily determine which percent that is. Since their bodies heal very easily, their ability to control their lives is promoted best if 99 percent of each body is removed in such a way that these abnormal individuals end up with what are, for us, normal human bodies. (Lippert-Rasmussen, 2008: 109).

Thought experiments like ‘Two hundred legs and arms’ are designed to capture the morally relevant features of a discussion, in this case the discussion of self-ownership, in order to make a specific point. The idea is to isolate the morally relevant features while controlling for all sorts of other features that may distort our moral judgments about the case in question. I have tried to make due without extremely outlandish thought experiments in this thesis, since I know that some readers may find them off-putting. I have refrained from using such thought experiments for practical purposes, and not because I believe they have no justified role to play in the method of analytic political theory.

As an example of what I believe is a relatively well-constructed thought experiment from this thesis, consider the thought experiment ‘Wiretapping’ from *Too Much Info*:

Wiretapping

Smith and Jones are neighbors. Unbeknownst to Jones, Smith wiretaps Jones’ telephone, using a fancy device which allows Smith to listen in on Jones’ conversations without violating Jones’ property rights. As it happens, Jones is on vacation for several months, and does therefore not use the telephone in that time period. (Mainz & Uhrenfeldt, 2021).

Some might say that the ‘fancy device’ referred to in this thought experiment makes it somewhat outlandish. But what we had in mind is not some sort of magic device,

but rather something like a tiny wiretap that can be placed in a way that does not violate Jones' property rights. What the thought experiment is meant to show, is that it is possible to violate the right to privacy without actually accessing any personal information about other people. The end goal of *Too Much Info* is to show that the Control Account of the right to privacy is superior to the Access Account. According to the latter, it is impossible for the right to privacy to be violated without someone getting access to personal information about other people. But Wiretapping shows that Smith does not get access to any personal information about Jones, and yet it seems plausible that Smith violates Jones' right to privacy. Further, it seems that the reason why Smith violates Jones right to privacy, is that Jones is not in control over whether or not Smith accesses the information. It was due only to contingent circumstances (that Jones happened to be on vacation) that Smith did not gain access. Thus, Wiretapping is designed to isolate the morally relevant features of 1) control, and 2) access, and demonstrates that - pace the Access Account - the right to privacy can be violated when neither 1) nor 2) is present.

2.4. Reductio ad Absurdum

The final methodological component that I want to describe is that of Reductio ad Absurdum. Suppose that we have a normative premise in an argument. It could, for instance, be a moral principle P that we believe to be true. The method of Reductio ad Absurdum then consists in searching through the logical implications of P, to see if

any of the implications are false. We might, for example, use our logical inference rules to find out what P implies. Suppose that P implies Q. Then we carefully consider the truth-value of Q. If Q is false, then it follows that P must be false as well. Now we know that we have a false premise in our argument, and it thus cannot be sound. The logical structure of Reductio ad Absurdum is the following, and it is called Modus Tollens:

Premise 1) $P \supset Q$

Premise 2) $\neg Q$

Conclusion: $\neg P$

It simply says that if Q follows from P, but Q is false, then P is false as well. In a previous subsection, we saw how we normally make use of moral judgments, and how we use thought experiments to pump these moral judgments. Those thought experiments are often supposed to establish the truth of Premise 2 in a Modus Tollens argument, namely that the implication of P is false. If we have a series of thought experiments where it seems obvious that $\neg Q$, and Q follows from P, then we have a good reason to think that $\neg P$.

Let me close this methodology chapter with an example of how I have used Reductio ad Absurdum in one of my papers. The example is from *An Indirect*

Argument. The reductio is meant to show that the version of the Control Theory defended by Leonhard Menges is mistaken. According to Menges, control theorists should interpret control as what he calls ‘source control’. He writes:

My main proposal is that privacy theorists can and should spell out privacy in terms of source control. According to the resulting source control account of privacy, an agent has privacy with regard to a certain piece of information just in case the person is the right kind of source of the relevant information flow if the information flows at all. In other words: 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. (Menges, Forthcoming: 9).

Menges leaves it underspecified what exactly it means to be the right kind of source of an information flow. However, regardless of what it means exactly, the idea of being the right kind of source of an information flow seems to be orthogonal to the issue of having privacy. I will return to this later, but for now, it suffices to illustrate how I use Reductio ad Absurdum to show that Menges’ view is mistaken. On page 17 in *An Indirect Argument*, I present the following (admittedly somewhat outlandish) hypothetical:

Moving Day

Every citizen of Private Ville lives in regular houses made of bricks. Every citizen of Private Ville is being wiretapped against his or her will by someone from outside of Private Ville. One day, every citizen of Private Ville chooses to move to houses that are made of fully transparent glass. Everyone that walks by such a house can see everything that happens inside the house. And, because the walls are made of thin glass, everyone outside the house can also hear every little sound from inside the house. No one is wiretapping the citizens of Private Ville in the new houses. But the people who were doing the wiretapping, are now standing outside the glass houses, watching and listening to what citizens of Private Ville do inside their houses. The citizens of Private Ville are fully aware of this. (Mainz, 2021b).

On the assumption that the citizens of Private Ville exercise source control when they choose to live in transparent houses, it follows that the citizens of Private Ville have full privacy. However, this seems very odd. The information about what the citizens of Private Ville do inside of their homes used to flow to one set of individuals. But by moving into the transparent house, the information now flows to the same set of individuals, but also to an additional set of individuals. To wit, more people now have access to their personal information, and those that have access have access to *a lot* of information. And yet, it follows from Menges' source control view that by moving

into the transparent houses, the citizens of Private Ville are performing a privacy *enhancing* act. The reductio thus consists in showing that Menges' theory (P) has the implication that the citizens of Private Ville have full privacy when they move into the transparent houses (Q). But, this implication is false ($\neg Q$), so it follows that Menges' theory is false as well ($\neg P$).

CHAPTER 3. What is This Thing Called Privacy?

In this chapter, I shall briefly prepare the theoretical ground for the papers to follow. As mentioned earlier, there are many competing theories of what privacy is, and perhaps especially, what the right to privacy is (assuming that this right even exists). Let me begin with a few disclaimers. Obviously, I cannot cover all aspects of privacy in this thesis. Not even close. There are many important and interesting topics in the privacy literature, and in the discussions of ethics of data analytics in general. Many of these topics are partly or completely left out in this thesis. For example, in the context of privacy, I only briefly discuss the non-informational aspects of privacy, such as ‘decisional privacy’. I do not discuss the important and interesting feminist theories of privacy. In the context of ethics of data analytics more broadly, I leave out important and interesting topics such as algorithmic fairness⁶, accountability⁷, transparency⁸, trust⁹, and so on. This delimitation is not motivated by any particular

⁶ See e.g. (Hedden, 2021); (Hellman, 2020); (Barocas & Selbst, 2016); (Binns, 2018); (Pessach & Shmueli, 2020).

⁷ See e.g. (Castets-Renard, 2019).

⁸ See e.g. (Blacklaws, 2018).

⁹ See e.g. (Kim & Routledge, 2020).

view of which topics are more important. Rather, it is motivated partly by relevance-considerations, and partly by interest-considerations.

Throughout the thesis, I assume that privacy rights exist. However, I am completely open to the view that there exist privacy wrongs that do not amount to violations of privacy rights. I am also open to the idea that privacy-related wrongs are best accounted for in consequentialist terms. I simply use the language of rights because it is the most common one in the literature. When I assume that privacy rights exist, I do not commit to any particular view of rights in general. I shall say a bit more about this in some of the papers to follow, but I generally strive to remain non-committal on the questions of the strengths of privacy rights, and how they fit in with more general conceptions of rights.

My co-supervisor Frej Klem Thomsen told me that Kasper Lippert-Rasmussen once suggested to him that philosophers need only be consistent within the bounds of any one paper. Following this (probably half-joking) advice surely makes it a whole lot easier to write a PhD thesis while being on a three-year long learning curve. However, I have of course tried to avoid making contradictions between two or more of the papers to follow. I probably did not succeed completely in avoiding contradictions. I am especially worried about certain tensions between the conclusions defended in *Too Much Info* and *Reply to Munch and Lundgren* on the one hand, and the conclusion defended in *Inferences*. I leave it for another occasion to

explore whether these tensions are problematic, and if so, what can be done to plausibly resolve the tension. But first, let us turn to the mess of defining privacy.

3.1. The Mess of Defining Privacy

As anyone who has been diving into the privacy literature knows; it is a mess. The only consensus in the literature seems to be that there is little consensus about anything else. There is no consensus about what the concept of privacy is. There is no consensus about what the right to privacy is, and there is no consensus on what the right to privacy, if it even exists, is there to protect (Marmor, 2015). The difficulty of defining privacy has led some theorists to think that there is no workable definition of privacy, but rather multiple types of privacy related to each other only by loose family resemblances (Solove, 2008). Some of the papers and chapters of this thesis attempt to show that while many things conventionally referred to as ‘privacy’ are at most loosely connected, there is still something to be said about how to conceptualize privacy. This thesis thus serves the role of clarifying some of the misconceptions and misunderstandings in the literature, but it would be naive to think that it by any means cleans up the mess entirely.

Many of the theoretical discussions of privacy did not - and does not - take place in philosophy outlets. A significant part of the discussion takes place in law journals.

It is often claimed that the academic discussion of privacy started with a paper in Harvard Law Review by Louis Brandeis and Samuel Warren called ‘The Right to Privacy’ (Warren & Brandeis, 1890). They argued that the right to privacy is ‘the right to be let alone’. The recent popularization of photography and newspapers had spawned the idea that people’s privacy needed to be better protected by law. Their paper received a lot of attention, and soon courts started to acknowledge the right to privacy. However, it was not until much later that the right to privacy was described and expanded more systematically. A new systematic treatment of the right to privacy was carried out by William Prosser in 1960. He introduced what he saw as the four essential interests in privacy:

1. Intrusion upon a person’s seclusion or solitude, or into his private affairs.
2. Public disclosure of embarrassing private facts about an individual.
3. Publicity placing one in a false light in the public eye.
4. Appropriation of one’s likeness for the advantage of another (Prosser, 1960: 389).

Both Warren and Brandeis, and Prosser, were concerned with moral problems that may arise when people gain access to personal information about others. A few years later, a different type of privacy rights began to be recognized in the law. This new

type of privacy right, which has been called ‘decisional privacy’, is a constitutional right to privacy, which was recognized by the Supreme Court. The constitutional right to privacy was announced in *Griswold v. Connecticut* (381 U.S. 479) in which the conviction of the Director of Planned Parenthood and a doctor from Yale University for distributing information about contraceptive to married couples, was overturned. The constitutional right to privacy was later applied to another famous case, namely *Roe v. Wade* (410 U.S. 113) involving the right to abortion. Suddenly, the right to privacy did not only cover informational privacy, but also a right to make certain ‘private’ decisions oneself without the interference of others, especially the state. Although decisional privacy so construed is very important, many theorists have pointed out that the very notion of decisional privacy confuses privacy with liberty.

After these initial advances in the privacy literature, it was not until much later that new philosophically interesting developments occurred. In 1975, Judith Jarvis Thomson published a paper in *Philosophy and Public Affairs* with the same title as Warren and Brandeis’ paper: ‘The Right to Privacy’ (Thomson, 1975). The publication of Thomson’s paper was the starting point of several of the major contemporary discussions in the philosophical literature on privacy. For example, Thomson’s paper spawned much of the contemporary discussion about ‘privacy

reductionism’.¹⁰ Moreover, a significant part of the contemporary discussion about the Control Theory and the Access Theory – which several of the papers included in this thesis is concerned with – was also spawned by Thomson’s paper.

Many privacy theories – especially the Control Theory, and the Access Theory - come in many different versions, and many of them come in both descriptive- and normative versions (Moore, 2008). The descriptive theories often try to explain what the *concept* of privacy is, or what it means to *have* privacy, or what it means to be in a *condition* of privacy. These theories are morally neutral. They do not, in themselves, imply what if anything is morally objectionable about privacy diminishments when they are. Neither are they ‘moralized’ in the sense that they take privacy to be an inherently normative term. The normative theories, on the other hand, often try to explain what the *right* to privacy is, which types of actions count as *violations* of this right, or what is morally objectionable about privacy diminishments when they are. Theories that hold that privacy is an inherently normative concept are – unsurprisingly – also conventionally categorized as normative privacy theories.

¹⁰ For an example of an earlier privacy reductionist, see (Davis, 1959).

Some theorists think that the distinction between descriptive- and normative theories is mistaken or misleading (Lundgren, 2020), or that the distinction does not exist because privacy is inherently a moralized concept (Inness, 1992). The account of privacy I defend in this thesis, is a non-moralized account of privacy. The primary reason for defining privacy in non-moralized terms is that I think it is perfectly sensible to talk about *having* privacy in a normatively neutral way. It is perfectly sensible to say that I can lose privacy without implying that anyone has done anything wrong, or that anyone's interests have been bettered or worsened. This does not preclude me from talking about privacy *rights*, or morally objectionable privacy diminishments, we just need different theories for the normative aspects of privacy.

Surprisingly, there has been little discussion about what the relation is between the concept of privacy, and the right to privacy. In fact, it is often frustratingly difficult to find out which theories are morally neutral, and which are not. In 2020, Björn Lundgren published a paper in *The Journal of Ethics*, in which he defended a particular view of what the relation is between the concept of privacy, and the right to privacy (Lundgren, 2020). As we shall see in *Reply to Munch and Lundgren*, I do not think that Lundgren gets this relation completely right.

An example of a descriptive privacy theory is the one defended by William Parent:

[p]rivacy is the *condition* of not having undocumented personal knowledge about one possessed by others. (Parent, 1983: 269).

Parent's theory is descriptive, because it is concerned with what it means to be in a *condition* of privacy. For Parent, having privacy is a question of being in a condition where others do not possess personal information about you. Similarly, Jeffrey Reiman suggests that

[...] privacy is the condition in which others are deprived of access to you (Reiman, 1995: 30).

Parent and Reiman's definitions are *descriptive* versions the *Access Theory* of privacy. They are concerned with what it means to *have* privacy, and they hold that having privacy is essentially a function of other people's access to certain personal matters. I shall return to the Access Theory shortly.

An example of a normative privacy theory is the one defended by Adam Moore:

Definition: A right to privacy is a right to control access to and uses of—places, bodies, and personal information (Moore, 2008: 421).

Moore's definition is explicitly concerned with the right to privacy, and suggests that this right consists in controlling the access to certain things. Note that the word

‘access’ is a part of Moore’s definition. Even so, Moore’s theory is a *normative* version of the *Control Theory*. The reason is that some of the normative versions of the Control Theory – like Moore’s – hold that the right to privacy is not a right to control things like personal information as such, but rather a right to control the *access* to the information. Things are already becoming complicated, so in the rest of this section I shall try to clarify what some of the different versions of the Control Theory and the Access Theory, respectively, hold. There are many privacy theories that do not fit nicely under the rubrics of either the Control Theory or the Access Theory. However, for present purposes it is useful to label many of the influential privacy theories as instances of either the Control Theory or the Access Theory. We begin with the former.

Control Theory

The Control Theory (sometimes referred to as the ‘control account’) is probably *the* most popular privacy theory.¹¹ As mentioned, the Control Theory comes in

¹¹ Different versions of the Control Theory can be found in (Westin, 1970); (Fried, 1968); (Moore, 2008); (Moore, 2010); (Parker, 1974); (Parent, 1983); (Allen, 2003); (Roessler, 2005); (Benzanson, 1991); (Goldberg, Hill, & Shostack, 2001); (Altman, 1976); (Calo, 2011); (Margulis, 1977); (Scanlon, 1975); (Beardsley, 1971); (Inness, 1992); (Menges, Forthcoming); (Mainz & Uhrenfeldt, 2021).

descriptive versions, and in normative versions. According to the descriptive versions, having privacy is essentially about having control over certain personal objects, such as personal information. Charles Fried writes:

As a first approximation, privacy seems to be related to secrecy, to limiting the knowledge of others about oneself. This notion must be refined. It is not true, for instance, that the less that is known about us the more privacy we have. Privacy is not simply an absence of information about us in the minds of others; rather it is the control we have over information about ourselves. (Fried, 1968: 484).

According to Fried, one's privacy is not diminished just because others have information about one in their minds. Rather, one's privacy is only diminished if one does not have control over *whether or not* they have this information.

Fried thinks that descriptive Access Theories like Reiman's or Parent's – which we shall return to shortly - are mistaken because they have implausible implications. As mentioned earlier, Fried uses the examples of a man on a desert island, and claims that

to refer [...] to the privacy of a lonely man on a desert island would be to engage in irony. (Fried, 1984: 209-210).

According to Fried, it does not make sense to say that a man stranded on a desert island enjoys privacy, because “the person who enjoys privacy is able to grant or deny access to others” (Fried, 1984: 210).

A descriptive Control Theory similar to that of Fried, comes from Richard Parker:

Privacy is control over when and by whom the various parts of us can be sensed by others. By "sensed," is meant simply seen, heard, touched, smelled, or tasted. By "parts of us," is meant the part of our bodies, our voices, and the products of our bodies. "Parts of us" also includes objects very closely associated with us. By "closely associated" is meant primarily what is spatially associated. The objects which are "parts of us" are objects we usually keep with us or locked up in a place accessible only to us. (Parker, 1974: 216).

A more recent descriptive Control Theory comes from Leonhard Menges. Menges' version of the Control Theory is motivated by numerous objections to the earlier versions of the Control Theory that target the unspecified use of the crucial word 'control'. According to Menges, control theorists should interpret control as what he calls 'source control'. This notion of control is inspired by the classic Frankfurt-cases

known from the literature on free will. Ironically, Menges leaves it very underspecified what he means by ‘source control’. He writes:

My main proposal is that privacy theorists can and should spell out privacy in terms of source control. According to the resulting source control account of privacy, an agent has privacy with regard to a certain piece of information just in case the person is the right kind of source of the relevant information flow if the information flows at all. In other words: 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. (Menges, Forthcoming: 9).

In *An Indirect Argument*, I offer a critique of Menges’ descriptive version of the Control Theory, and show how it collapses into a descriptive version of the Access Theory.

The descriptive versions of the Control Theory have been criticized by many, especially access theorists. An influential critique comes from William Parent:

All of these definitions [the Control Theories, red.] should be jettisoned. To see why, consider the example of a person who voluntarily divulges all sorts of intimate, personal, and undocumented information about himself to a friend. She is doubtless exercising control. . . . But we would not and should

not say that in doing so she is preserving or protecting her privacy. On the contrary, she is voluntarily relinquishing much of her privacy. People can and do choose to give up privacy for many reasons. An adequate conception of privacy must allow for this fact. Control definitions do not. (Parent, 1983: 273).

Many – myself included - have found Parent’s counterexample convincing. As will become clear in *An Indirect Argument*, I think that there are simply too many, and too convincing, counterexamples to the *descriptive* versions of the Control Theory for it to be plausible. Instead, as I shall explain later, I prefer what Lauritz Munch and I have called the Hybrid View of privacy. But let us not get ahead of ourselves. We still need to be acquainted with some of the *normative* versions of the Control Theory.

We have already seen one of the normative version of the Control Theory, namely the one defended by Moore. A more recent one comes from Andrei Marmor:

... a general right to privacy [is] grounded in people’s interest in having a reasonable measure of control over the ways in which they can present themselves (and what is theirs) to others. (Marmor, 2015: 3-4).

It should be fairly easy to see how Marmor’s account counts as a normative version of the Control Theory, despite the fact that it does not say explicitly what this right *is*,

but rather what it is ‘grounded in’, namely an interest in a reasonable amount of *control* over how they can present themselves to others.

A lot can be said about each of the theories mentioned above, and I have done very little to spell out how we should understand each of them. But the primary purpose has been to show some of the various versions of the Control Theories, and especially to show that they can come in descriptive- as well as normative versions. In *Too Much Info*, and in *Reply to Munch and Lundgren*, Rasmus Uhrenfeldt and I develop a normative version of the Control Theory that is meant to take into account many of the objections that have been raised to older versions of the theory, like the ones mentioned above. Next, let us turn to some of the different versions of the Access Theory that can be found in the literature.

Access Theory

The Access Theory (also often referred to as the ‘access account’) has historically been less popular than the Control Theory. However, very recent contributions to the literature, such as my paper *An Indirect Argument*, and the papers from theorists like

Björn Lundgren (2020) and Kevin Macnish (2018) show that the Access Theory is alive and kicking.¹²

Just like the Control Theory, the Access Theory comes in both descriptive and normative versions. According to the descriptive versions, one's privacy is diminished to the extent that others have access to one's personal objects, such as one's personal information. Some versions of the Access Theory allows for privacy being a matter of degree, so that one can have more or less privacy, depending on the number of people who have access to the private objects, and the number of objects they have access to. As will become clear later, I favor a version of the descriptive Access Theory that allows for privacy being a matter of degree along these two dimensions, plus a third dimension: the dimension of the strength of epistemic relations. To wit, one's privacy over one's personal objects is a function of a) the number of people who have access to the personal objects, b) the number of personal objects they have access to, and c) the strength of the epistemic relation that they have to the personal objects. Classic versions of the Access Theory do not allow for privacy being a matter of degree as a function of c).

¹² Different versions of the Access Theory can be found in (Thomson, 1975); (Gavison, 1980); (Bok, 1989); (van den Haag, 1971); (Macnish, 2018); (Lundgren, 2020); (Mainz, 2021b).

Let us look at an example of a descriptive version of the Access Theory. This version has been defended by Ruth Gavison. She writes:

Our interest in privacy, I argue, is related to our concern over our accessibility to others: 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. (Gavison, 1980: 423).

According to Gavison, privacy is about accessibility to others. This is in clear contrast to the claims from the control theorists we saw in the previous section, especially the one from Charles Fried. A more recent defense of the Access Theory comes from Kevin Macnish. He provides the following thought experiment to vindicate the conclusion that the Access Theory is correct:

Despite its relative unpopularity, I believe that the access account is correct. This can be illustrated through returning to the diary example. Imagine that I have returned to the coffee shop after a 30 minute interval to find my diary on a stranger's table. It is unopened. 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.

I feel an enormous sense of relief, thank her and leave with my dignity intact.

(Macnish, 2018: 420-421).

Macnish thinks that privacy has not been *lessened*, nor *violated* in the diary example. But clearly there is a difference between a privacy loss and a privacy violation. One can certainly lose privacy without having one's right to privacy violated. The access theorists need to accept this too. It would be a devastating *reductio* to their own view, if they maintained that person X getting access to private object Y about person Z is a sufficient condition for X violating Z's right to privacy. If, for example, Z chose, voluntarily, to give X access to Y, it would be very strange to say that Z has now violated X's right to privacy. Thus, it seems necessary for the access theorists to conceptually separate the descriptive and the normative versions of the Access Theory.

Let us now turn to the normative versions. Normative versions of the Access Theory typically hold that a privacy diminishment is a necessary condition for the right to privacy to be violated.¹³ But for reasons spelled out above, they do not hold that privacy diminishments are sufficient for the right to privacy to be violated. In this

¹³ For an example of a theorist who explicitly denies this, see (Kappel, 2013).

sense – and this is probably controversial - normative versions of the Access Theory can be seen as *adding* a necessary condition to the normative versions of the Control Theory, namely that someone actually access the personal information in question. This view is most famously defended by Judith Jarvis Thomson. She writes:

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. (Thomson, 1975: 304).

Thomson's thinks that normative Control Theories are mistaken, because she thinks that it is possible to lose control without the right to privacy being violated. The fact that your neighbor invents an X-ray means – Thomson thinks – that you lose privacy over the access to your personal information. But surely, your right to privacy is not violated until your neighbor actually trains the X-ray on the wall.

In *Too Much Info*, Rasmus Uhrenfeldt and I argue that control theorists need not interpreted 'control' in the way that Thomson does, and consequently, they can avoid Thomson's objection. We use this to show that Thomson's normative version of the Access Theory is mistaken. However, in *Reply to Munch and Lundgren*, we

backtrack on this idea, and instead show that a Thomson-like normative version of the Access Theory is compatible with a specific normative version of the Control Theory.

3.2. The Concept of Privacy and the Right to Privacy

Much of the confusion in the privacy literature stems from the fact that many privacy theorists do not clarify whether they are talking about the concept of privacy, or the right to privacy, or something else entirely. This confusion has been evident in the privacy literature more or less throughout the entirety of its history. William Parent made a similar point already in 1983 (Parent, 1983: 269). However, like many other discussions in the privacy literature, there is no consensus about what the relation is between the concept of privacy, and the right to privacy. In this section, I shall briefly sketch out one way to think of this relation that I find particularly helpful. I return to this in *Reply to Munch and Lundgren*.

The basic idea is that it is one thing to be in a condition of *having* privacy, and another thing to have a *right* to privacy. You can lose privacy, without someone violating your right to privacy, but your right to privacy cannot be violated without someone diminishing your privacy. To wit, privacy diminishments are necessary but not sufficient for privacy violations, and privacy violations are sufficient but not necessary for privacy diminishments. The distinction between the concept of privacy

and the right to privacy is already familiar from the previous section on descriptive- and normative privacy theories. The descriptive ones are generally concerned with the concept of privacy, and the normative ones are generally concerned with the right to privacy.

To see why a privacy diminishment is not sufficient for a privacy violation, recall Parent's objection to the Control Theory:

Parent's Objection

All of these definitions [the control definitions of privacy] should be jettisoned. To see why, consider the example of a person who voluntarily divulges all sorts of intimate, personal, and undocumented information about himself to a friend. She is doubtless exercising control, in a paradigm sense of the term, over personal information about herself as well as over (cognitive) access to herself. But we would not and should not say that in doing so she is preserving or protecting her privacy. On the contrary, she is voluntarily relinquishing much of her privacy. People can and do choose to give up privacy for many reasons. An adequate conception of privacy must allow for this fact. Control definitions do not. (Parent 1983: 273).

Parent's Objection is an objection to the Control Theory because it shows that one can lose privacy without losing control. Parent's Objection describes a scenario where

privacy is supposedly lost even though control is not lost. The individual in Parent's Objection chooses to divulge all sorts of personal information to a friend. The individual has clearly lost privacy relative to that friend and to the personal information in question, but the individual was in control at all times during the process of giving the friend access to the information. According to Parent, the Control Theory cannot explain why the person loses privacy, and yet it is clear that she does.¹⁴

As Björn Lundgren has recently pointed out, control theorists often give a specific reply to Parent's Objection. I denote it the 'Standard Reply', and it can be paraphrased like this:

Standard Reply

Parent's counterexample misses its target, because the individual who chooses to divulge personal information to a friend does in fact lose control, because it is not within that individual's control whether the friend will pass

¹⁴ More recently, Leonhard Menges has called objections like Parent's 'voluntary divulgence objections' (Menges, Forthcoming).

the information on to others (Gavison, 1980: 427); (Lundgren, 2020: 168-169); (Matheson 2007: 255).

According to the Standard Reply, the individual in Parent's Objection does in fact lose control, since she is no longer in control over whether the friend will pass on the information to others. As will become clear in *An Indirect Argument*, I am not convinced by the Standard Reply. Any plausible privacy theory must be consistent with the verdict that the individual who voluntarily chooses to divulge personal information to a friend in fact loses privacy with regards to that friend and those pieces of information. But, any plausible privacy theory should also hold that voluntarily divulging information is not sufficient for one's right to privacy to be violated (Fallis, 2013). So, intuitively, losing privacy, or having one's condition of privacy 'reduced' is not a sufficient condition for a violation of the right to privacy to occur.

To see why a privacy diminishment is necessary for a privacy violation (and thus why a privacy violation is sufficient for a privacy diminishment), consider the following. The right to privacy must be a right to *privacy*, as opposed to a right to something else. But if a violation of the right to privacy can occur even without any change to the right-holder's condition of privacy, then it is difficult to see how the right to privacy is a right to *privacy* in particular (Lundgren, 2020). In other words, if a privacy diminishment is not necessary for a violation of the right to privacy, then

the relation between being in a condition of privacy, and having a right to privacy, appears to be rather puzzling.¹⁵

The consideration that the right to privacy must be a right to *privacy* as opposed to a right to something else, has led some theorists to think that it is conceptually impossible to subscribe to one type of theory at the ‘descriptive level’, and another type of theory at the ‘normative level’ (Lundgren, 2020). They seem to think that whatever the definiens of the descriptive condition of privacy is, the right to privacy must have the same definiens. The idea is that because there must be a close relation between the concept of privacy and the right to privacy, one cannot – for instance – subscribe to an Access Theory of the concept of privacy, and a Control Theory of the right to privacy. However, I believe that this view is mistaken. In fact, my view is exactly that the concept of privacy should be defined in terms of access, and that the right to privacy should be defined in terms of control. In a later section, I will briefly defend the view that the concept of privacy is best conceived of as particular version of the Access Theory, which Lauritz Munch and I call the Hybrid

¹⁵ Despite all of this, some theorists do in fact believe that privacy diminshments are not necessary for privacy violations. See for instance (Kappel, 2013); (Skopek, 2020).

View of privacy. And, as will become clear in *Too Much Info* and *Reply to Munch and Lundgren*, I defend a certain version of the Control Theory of the right to privacy.

So, what is the relation between the concept of privacy, and the right to privacy? As Rasmus Uhrenfeldt and I explain in *Reply to Munch and Lundgren*, the concept of privacy should be defined in terms of access, but the right to privacy so construed is a control right similar to the type of control right that are conventionally thought to be a part of the bundle of rights that constitute property rights. On this view, the Control Theory explains the type of control right that one has over the access to one's personal information. And, if this right is violated, then one's privacy is diminished because someone accesses one's personal information. Perhaps all of this is still very abstract, and difficult to wrap one's head around. Hopefully, it all becomes more clear in the papers to follow.

3.3. The Wrongness of Privacy Violations

A relatively large part of this thesis is concerned with questions related to privacy violations. However, I do not discuss in any of the papers what grounds the right to privacy. I do, however, discuss the question of which types of actions constitute violations of the right to privacy. In an earlier section, I introduced a range of research questions. To get a better grasp of what the respective aims of my papers are, and what

they are not, it is helpful to point out the differences between the following three questions:

Question 1: How should ‘privacy’ be conceptualized?

This question asks – at a basic level - what privacy is. Answering this question does not necessarily commit one to a particular view of what the right to privacy is, or what makes privacy reductions wrong when they are. Descriptive privacy theories – like the ones we saw in earlier sections - normally seek to answer this question. *Reply to Munch and Lundgren*, and *An Indirect Argument*, both offer partial replies to Question 1. So does the section of this thesis called ‘The Hybrid View of Privacy’.

Question 2: Which types of actions constitute violations of ‘the right to privacy’?

This question asks which actions count as violations of the right to privacy, and which do not. Theories that try to answer Question 2 do not necessarily try to explain *why* certain actions constitute violations of the right to privacy. For instance, not all normative versions of the Control Theory and the Access Theory described in earlier sections try to explain what grounds the right to privacy, but they do try to distinguish privacy violations from non-violations. *Too Much Info*, *Reply to Munch and Lundgren*, and *Inferences* offer partial replies to Question 2.

Question 3: What grounds ‘the right to privacy’?

This question asks *why* an action that constitutes a violation of the right to privacy is wrong, or – perhaps more broadly – what interests the right to privacy protects. None of the papers included in this thesis attempt to answer Question 3, so I shall shortly sketch out some of the attempts to answer it that can be found in the literature.

One might wonder if any correct answer to Question 3 implies an answer to Question 2, and vice versa. If so, then perhaps we do not need an answer to both of these questions. This would be good news for me, since I do not explicitly discuss Question 3 in any of my papers, while I do discuss Question 2. Plausibly, a correct answer to Question 2 is likely to have *some* implications for the answer to Question 3, and vice versa. But there are at least two reasons for thinking that it does not suffice to only address one of these two questions, if one wants to get a good grip of the normative aspects of privacy. The first reason is that even if we, for example, have a good answer to Question 2, there may well be several competing answers to Question 3. We may be left with a plurality of accounts, each of which explains why a certain

action is a violation of the right to privacy.¹⁶ The second reason is that there may be certain clashes or incoherences between answers to Question 2 and answers to Question 3. This is one reason why we use the method of Reflective Equilibrium to try to balance our judgments about whether particular types of actions count as violations of the right to privacy, with more general accounts of what makes said actions objectionable.

Before turning to two competing accounts of what grounds the right to privacy, let me briefly set Question 1-3 aside from some of the other privacy-related issues covered in the rest of the papers included in this thesis. *Reply to Cheneval* does not directly involve Question 1-3. At least, the paper does not try to answer any of these questions directly. However, the idea of granting individuals property rights over personal data is often motivated by an attempt to protect individuals' privacy in the sense discussed in *Too Much Info*, *Reply to Munch and Lundgren*, and *An Indirect Argument*. Thus, if all I say in *Reply to Cheneval* is correct, then granting people legal property rights over their personal information might not be a way to protect their

¹⁶ The view that there are indeed a plurality of interests that are harmed by privacy violations seems to be widespread in the literature. See e.g. (Kappel, 2013); (Rachels, 1975).

moral property rights, but rather their moral privacy rights. However, if what I say in *Reply to Cheneval* is mistaken, then it may have implications for how Question 1-3 are best answered. It might be the case that what we normally call ‘privacy’, and what we normally call ‘privacy rights’ are empty notions that perhaps ought to be replaced with notions related to ownership and property rights instead.

In *How to Buy an Election*, we discuss some of the normatively important consequences of losing privacy in the sense described in *An Indirect Argument*, and in the section of this thesis called the Hybrid View of Privacy. However, the argument in *How to Buy an Election* does not concern privacy *violations*. That is, it is not trying to answer Question 2, nor Question 3. Rather, it is concerned with discussing non-privacy related normative questions that arise when people lose privacy in a non-normative sense. The normative question discussed in *How to Buy an Election* is one that arises when it is possible to use data analytics to predict certain pieces of personal information about individuals, namely how they will vote in an upcoming election. However, the normative question that arises is not one that revolves around the *right* to privacy, but rather one that revolves around the presumed problem of buying an election outcome.

Let us now turn to two competing accounts of what grounds the right to privacy. That is, the two following accounts attempt to answer Question 3, and as mentioned, they may have some bearing on how best to answer Question 1 or

Question 2. Following (Munch, 2021), I shall call them ‘the autonomy account’, and the ‘subsequent harm account’, respectively.

The Autonomy Account

The autonomy account holds, roughly, that privacy rights are essentially grounded in the value of autonomy. That is, privacy is important because it somehow protects individuals’ autonomy. This account has been discussed by prominent theorists, including James Stacey Taylor (2002) and Joel Feinberg (1986).¹⁷ What exactly the autonomy account holds depends in turn on which underlying account of autonomy we adopt. One oft-cited account of autonomy is that of Joseph Raz. According to Raz, for an agent A to be in a state of autonomy, A must have - *inter alia* - an adequate range of valuable options available to her to choose from, and the relevant independence to choose freely among those options (Raz, 1988).

Assuming Raz’ account of autonomy, the resulting autonomy account of privacy holds that privacy violations are objectionable because they either leave A with an inadequate set of valuable options to choose from, or somehow renders A

¹⁷ See also (Mokrosinska, 2018); (Henkin, 1974); (Lippke, 1989) for more discussions of privacy and autonomy.

unable to independently choose from the set of options (by coercion, manipulation, or something similar) (Munch, 2021). Suppose, for instance, that A has a legitimate interest in keeping her medical record to herself. Now suppose that B, a very skilled hacker, gains access to A's online copy of her medical record. In this examples, A does not have autonomy in any relevant sense, because she does not have an adequate set of valuable options available to her. Supposedly, A has no option but to live with the fact that B now knows what medical conditions she suffers from. Information that A wanted to keep to herself. Now suppose that B poses as a medical doctor and asks A to hand over the copy of her medical record. A thinks that B is a real medical doctor, but in fact he is not. We might plausibly say that A does not have autonomy in any relevant sense, because A is unable to independently choose from the relevant set of valuable options, because she is being manipulated.

The Subsequent Harm Account

According to the subsequent harm account, privacy rights are grounded in the consideration that A often has an instrumental reason for wanting that her privacy is not diminished by B. The instrumental reason is that B might use the relevant information about A in ways that harm A – or make A worse off in a morally relevant

way - down the line.¹⁸ To return to the example of A and her medical record, let us suppose that B works for an insurance company. B wants to know A's current medical condition, so that B can more accurately predict the likelihood of A suffering from diseases that are expensive to treat. B therefore decides to put his hacking skills to use, and gains access to A's online copy of her medical record. Based partly on the information in A's medical record, B trains a machine learning algorithm and learns that there is a high likelihood that within a few years, A will suffer from a disease that is very expensive to treat. Consequently, B decides to increase A's premium drastically, and A is no longer able to afford insurance. In this example, B plausibly violates A's right to privacy. On the subsequent harm account, what grounds A's right to privacy is the risk that if B gets access to A's personal information, then B might use the information in a way that causes A harm.

The autonomy account and the subsequent harm account are competing theories of how best to answer Question 3. As mentioned above, none of the arguments defended in my papers force me to subscribe to any of these competing accounts. That said, some of my arguments in *Too Much Info* and *Reply to Munch and*

¹⁸ For discussion of the subsequent harm account, see (Parent, 1983); (Marmor, 2015); (Munch, 2020); (Munch, 2021).

Lundgren may sit uncomfortably with the subsequent harm account, if the notion of likelihood involved in this account is interpreted strictly. What I mean by this is that if it must be very likely that B uses A's personal information f in a way that harms A in order for A to have a right to privacy over f , then my arguments in those two papers are probably too wide. I claim in those papers that it is sufficient for A's right to privacy to be violated that she does not have the relevant kind of control over the access to f . The relevant type of control is what I call 'Negative Control'. But if this is correct, then it seems that one has privacy rights even over pieces of information that are very unlikely – or even impossible - to be used in objectionable or harmful ways.

With these remarks in place, let us now turn to the question of how best to answer Question 1. As mentioned, both *Reply to Munch and Lundgren*, and *An Indirect Argument* give partial answers to Question 1. In the following section, I try to flesh out a plausible answer to Question 1 in more detail. The view I defend is based partly on the view defended in (Munch & Mainz, 2021). I shall argue that A has privacy if, and only if, B does not have warranted beliefs about A, and B does not perceive A (or her personal matters). Call it the 'Hybrid View' of privacy.

3.4. The Hybrid View of Privacy

Until recently, epistemologists have paid little to no attention to privacy. In fact, some epistemologists have explicitly argued that privacy (for the most part) falls outside the scope of epistemology. For example, Alvin Goldman writes:

Important as [privacy] is, it does not squarely fall into the domain of epistemology as I have delineated it, because epistemology focuses on the means to knowledge enhancement, whereas privacy studies focus on the means to knowledge curtailment (at least decreasing knowledge in the hands of the wrong people). For this reason, I shall not explore this topic. I do not belittle the importance of privacy as a moral issue; it simply falls, for the most part, outside the scope of epistemology. (Goldman, 1999: 173).

In 2007, David Matheson took a completely different approach, and tried to show that privacy is essentially an epistemic concept, because your privacy seems to be completely dependent on whether or not others *know* some personal facts about you (Matheson, 2007). Matheson provides a range of effective objections to the classic versions of the Control Theory and the Access Theory, and he tried to show how his own account effectively avoids their shortcomings. However, Matheson's account has its own shortcomings. In 2013, in a special issue in *Episteme*, a range of epistemologists criticized Matheson's account for, *inter alia*, not allowing for privacy

to be a matter of degree.¹⁹ The idea is that others do not need to *know* some personal fact about you in order for your privacy to be reduced. Weaker epistemic relations than knowledge suffice. The stronger the epistemic relation someone has with regards to a personal fact about you, the more your privacy is reduced with regards to that person and that fact. In this section, I explain how the critique of Matheson offered by the privacy epistemologists can straightforwardly be built into Matheson's account. After that, I briefly explain why there is good reason to think that the resulting account is mistaken, and I suggest an alternative account.

What does it mean when we say that someone's privacy is 'reduced' or 'diminished'? Numerous attempts to answer the flipside of that question - what does it mean to *have* privacy? - have been proposed in the philosophical literature. If we want to know what it means when we say that someone's privacy is reduced, it might be a good starting point to ask what it means to have privacy in the first place. Presumably, if you go from *having* privacy, to having less privacy, then your privacy has been *reduced*. Matheson defined and criticized the most prominent accounts of what it means to have privacy. The first one is a version of the – by now familiar - Control Theory:

¹⁹ See (Blaauw, 2013); (Kappel, 2013); (Fallis, 2013).

CT An individual A has informational privacy relative to another individual B and to a personal fact f about A if and only if A controls whether B knows f . (Matheson, 2007: 252).

According to this version of the Control Theory, the desideratum for you having privacy is that you have *control* over whether someone else knows a personal fact about you. If someone else's knowledge of a personal fact about you is completely dependent on your voluntary choice to grant that person access to the fact in question, then - on the Control Theory- you have informational privacy. Matheson offers an objection to the Control Theory:

Suppose again that an individual voluntarily discloses all sorts of personal facts about herself to a friend. But now suppose further that, first, this individual is such an intimidating character that the friend in question would never dream of passing the facts along to others without her express permission, and that, secondly, this individual has the unusual ability to interfere with the friend's memory in such a way as to remove the friend's memorial knowledge of the facts even after she has acquired it. Now we can say that the individual has relinquished her informational privacy with respect to the personal information and the friend; that is, relative to the friend and to the range of personal facts of which that information consists, the individual has no informational privacy. Yet in this case [...], because the

individual remains in complete control both of whether the friend continues to know the facts and of who, other than the friend, comes to know them, it will follow on CT that the individual has not voluntarily relinquished her informational privacy at all—whether relative to the friend and to the personal facts, or relative to others and to those facts—a very counterintuitive result. (Matheson, 2007: 255-256).²⁰

Matheson's objection to the Control Theory hits the nail on the head. If you voluntarily give someone access to all personal facts about you, and you simultaneously can control whether that someone gives other people access to the same fact, then on the Control Theory you still have complete informational privacy. But it seems very counterintuitive that your privacy has not been reduced in relation to the person who now knows all sorts of personal facts about you. This is a very powerful objection to the Control Theory. Any plausible descriptive theory of privacy

²⁰ Matheson's thought experiment is an extension of Parent's famous counterexample to the Control Theory (Parent, 1983: 273). Matheson's extends the example, so that it becomes immune to what I called the Standard Reply in an earlier section.

must imply that the person in Matheson's thought experiment has her privacy reduced relative to the friend.

The second account discussed by Matheson is the so-called 'Limited-Access Theory' (LAT):

LAT An individual A has informational privacy relative to another individual B and to a personal fact f about A if and only if there are extraordinary limitations on B's ability to know f .²¹ (Matheson, 2007: 253).

According to the LAT, the desideratum for you having privacy is that there are extraordinary limitations to someone else coming to know a personal fact about you. As long as there are extraordinary limitations on other people's ability to know a personal fact about you, your informational privacy is intact. An example of an extraordinary limitation could be a legal norm according to which "... unauthorized videotaping of an individual's activity in her own home is subject to prosecution."

²¹ Examples of proponents of LAT are (Allen, 1988); (Gavison, 1980); (Bok, 1989) and many others. In an earlier section, I described the literature as basically split between versions of the Control Theory and versions of the Access Theory. I introduced this broad distinction for the sake of simplicity. I think, however, that it is useful to think of the LAT and the subsequent NIT and BIT as versions of the more broad Access Theory, although the presence of a strong control right might also count as an 'extraordinary limitation on B's ability to access f '.

(Ibid). Matheson refers to the following quote from William Parent, as an effective counterexample to the LAT:

[B] taps [A]’s phone and overhears many of her conversations, including some of a very intimate nature. Official restraints have been imposed on [B]’s snooping, though. He must obtain permission from a judge before listening in on [A]. (Parent, 1983: 274).

According to the LAT, it is a sufficient condition for your informational privacy that there are extraordinary limitations on other people’s ability to know some fact about you. In Parent’s example, this condition is satisfied, and yet there is clearly a loss of informational privacy.

The third theory discussed by Matheson is the so-called ‘Narrow Ignorance Theory’ (NIT):

NIT An individual A has informational privacy relative to another individual B and to a personal fact f about A if and only if (1) f is undocumented and (2) B does not know f .²² (Matheson, 2007: 253).

According to the NIT, the desideratum for you having privacy is that the fact in question is undocumented, which means that it is not ‘belonging to the public record’²³, and that the fact is not known by anyone.²⁴ Matheson refers to the following quote from Judith Wagner DeCew as an effective objection against the NIT:

[D]uring former Massachusetts Representative Margaret Heckler’s divorce proceedings, her husband claimed that they had not had sexual relations in twenty years. Although this information was publicly available to reporters in the courtroom, it seems clear that the subsequent media coverage not only diminished Heckler’s privacy but also violated her right to privacy. (DeCew, 1997: 30).

²² An example of a proponent of the NIT is (Parent, 1983), and I take it to be an instance of the broader Access Theory.

²³ This is what gives rise to the ‘Narrow’ part of the name of the theory.

²⁴ This is what gives rise to the ‘Ignorance’ part of the name of the theory.

This example clearly shows - pace the NIT - that informational privacy can be lost even though the personal fact in question is ‘documented’, i.e. belongs to the public record. The fact that thousands of people got epistemic access to the fact through the media coverage clearly diminished Heckler’s privacy, even though the personal fact was already publicly available.

Finally, Matheson suggests his own theory, which he calls the ‘Broad Ignorance Theory’:

BIT An individual A has informational privacy relative to another individual B and to a personal fact f about A if and only if B does not know f . (Matheson, 2007: 259).

According to the BIT, the desideratum for you having privacy is simply that someone else does not have knowledge about a personal fact about you. The BIT elegantly avoids all of the shortcomings of the Control Theory, the LAT and the NIT; 1) it can explain why the person who voluntarily divulges personal information to a friend loses privacy, 2) it can explain why someone whose phone is wiretapped loses privacy, even though a warrant is needed before the wiretapping begins, and 3) it can explain why Heckler loses privacy, even though her personal fact was already publicly available.

However, the BIT is not without its own problems. Don Fallis, for one, suggests that having *knowledge* about some fact about someone else is not necessary for reducing privacy. According to Fallis, a weaker epistemic relation than knowledge is sufficient, such as justified true belief, or something similar. Fallis writes:

Suppose that Sam simply had a vivid dream about Norm having a tattoo on his butt. When he wakes up, he finds that he cannot help believing that Norm has a tattoo. But there is no connection between the fact that Norm has a tattoo and Sam's belief. In that case, it does not seem that Norm has lost his privacy about the tattoo with respect to Sam (cf. Matheson 2007: 264; Peels 2012). In order for Norm to lose his privacy about the tattoo with respect to Sam, Sam's belief needs to be hooked up to Norm's tattoo in some way. The plausible suggestion of the knowledge account is that Sam has to know about the tattoo in order for Norm to lose his privacy. However, suppose that Cliff is an unreliable testifier, but that he is telling the truth when he tells Sam that Norm has a tattoo on his butt. In that case, Sam does not know that Norm has a tattoo on his butt. But it still seems that Norm has lost his privacy about the tattoo with respect to Sam. (Fallis, 2013: 157).

Klemens Kappel agrees with Fallis that certain epistemic relations weaker than knowledge can be relevant for privacy. He writes:

Insofar as there are answers permitting some level of precision, it is difficult to believe that they could sustain categorical views. By categorical views I mean views asserting, for example, that the epistemic state of knowing, and only that, matters for privacy. According to such views, justified but yet not known true belief has no effect whatsoever on privacy diminishment. This is hard to believe. If knowledge affects privacy, as it surely does, then so do epistemic states that are in many respects just like knowledge, among them highly justified true beliefs. (Kappel, 2013: 188).

Martijn Blaauw agrees with Fallis and Kappel, and suggests that privacy comes in degrees. He writes:

Even if there is just one personal proposition that I would like to have privacy about with respect to just one individual, there can still be flexibility in the degree of privacy I have regarding this proposition. This dimension has to do with the type of epistemic relation the individual stands in, if any, vis-à-vis the personal proposition in question. (Blaauw, 2013: 171).

Blaauw lists a number of different relations one can have with respect to a given personal fact (Ibid):

- (1) A mere belief that P

- (2) A true belief that P
- (3) A justified true belief that P
- (4) A degettierized true belief that P
- (5) A rational true belief that P
- (6) A warranted true belief that P
- (7) Knowledge that P
- (8) Certainty that P

Although not all of the relations in (1)-(8) should count as *epistemic* relations ((1) can be false, and (2) can be a result of mere wishful thinking), the basic idea is simple: The degree to which someone's informational privacy is reduced, depends on the epistemic relation someone else has to the fact in question. The closer the relation is to (8), the more informational privacy is reduced. The BIT does not allow for this, since it only includes relation (7); knowledge. In other words, informational privacy is not a matter of degree on an epistemic scale, according to BIT. However, BIT is compatible with privacy being a matter of degree along at least two other dimensions: The number of *people* who knows a personal fact, and the number of *facts* they know.

The objections from the privacy epistemologists give us at least *prima facie* reason to reject BIT. It seems plausible that the strength of the epistemic relation affects the degree to which privacy is reduced. But we need not abandon BIT altogether, since we can incorporate the ideas that not only a lack of knowledge, but also a lack of weaker epistemic relations are necessary and sufficient for having privacy. Call the new theory BIT2:

BIT2 An individual A has full informational privacy relative to another individual B and to a personal fact *f* about A if and only if there is no epistemic relation between B and *f*. A's privacy is reduced relative to the strength of the epistemic relation between B and *f*.

According to the BIT2, privacy is a matter of degree. When there is *any* epistemic relation between B and a personal fact about A, there is a reduction of A's privacy. Moreover, the stronger the epistemic relation is between B and the personal fact about A, the more significant the reduction of A's privacy is.

Let me now briefly explain why BIT2 is mistaken – despite it being immune to the objections from the privacy epistemologists. The critique I offer, is a condensed version of the critique offered by Lauritz Munch and me in the article ‘To Believe, or Not to Believe, That is not the (Only) Question: A Hybrid View of Privacy.’ (Munch & Mainz, 2021).

According to BIT2, the fact that B has a warranted belief in f is both necessary and sufficient for A's privacy to be reduced. The degree to which A loses privacy depends, *inter alia*, on the strength of the epistemic relation that B has to f . However, BIT2 is mistaken, because there are other ways in which one can lose privacy. It is not necessary that others form a warranted belief in order for privacy to be reduced. In fact, it is not necessary that others form a belief at all. To see why, consider

Defeat. B X-rays A's safe and observes the nude photo of A contained in the safe. However, B firmly believes that he hallucinates, so he forms no belief on the matter. (Munch & Mainz, 2021: 7).

It seems intuitively true that A's privacy is reduced relative to B (and that B violates A's right to privacy). But since B forms no beliefs about A, let alone forms a warranted belief, it cannot be the case that a warranted belief must be formed in order for privacy to be reduced. As Munch and I suggest, there is a further necessary condition that must be satisfied, in order for A to have full privacy relative to B, and with regards to f : that B is not in a perceptual state regarding f . To wit, B perceiving f is sufficient for

A to lose privacy relative to B, and with regards to f .²⁵ Accordingly, A has privacy if, and only if, B does not have a warranted belief that f , and B does not perceive f . We call this view

The Hybrid View: Individual A has privacy regarding relevant information p and with respect to individual B *iff* B lacks epistemically warranted belief that p , *and* B is not in a perceptual state regarding p . (Munch & Mainz, 2021: 12).²⁶

The Hybrid View handles cases like Defeat elegantly. On the Hybrid View, A's privacy is reduced because B perceives A's nude photo. By contrast, privacy accounts like BIT2 – and others that focus exclusively on beliefs, cannot handle cases like Defeat. But there is also another reason to favor the Hybrid View over belief-based accounts, namely that the former squares better with many of our intuitions

²⁵ I will not get into what the term 'perception' covers here. But see (Munch & Mainz, 2021) for elaboration.

²⁶ Note that in BIT2, the relevant information is denoted ' f ', while in (Munch & Mainz, 2021) we denoted it ' p '.

concerning the right to privacy. In particular, the Hybrid View squares well with our normative intuitions in cases of ‘repeated violations’, such as the following:

Sequence. At t_1 , A X-ray’s B’s safe and learns all the details of its contents. Prior to t_1 , A had no warranted beliefs (including disbelief) about the contents of B’s safe. At t_2 , A is in a warranted belief-state regarding the contents of B’s safe, but does not engage in X-ray activity. At t_3 , A X-ray’s B’s safe again. (Munch & Mainz, 2021: 15).

It is a plausible assumption that a reduction of privacy is a necessary condition for a violation of the right to privacy.²⁷ Now, if we accept this assumption, then belief-based accounts cannot explain the violation of the right to privacy that occurs in Sequence. Plausibly, there is a violation of B’s right to privacy at t_1 , and again at t_3 . However, on belief-based accounts like BIT2, there is no reduction in privacy at t_3 , because no warranted beliefs are formed at t_3 . Thus, if belief-based accounts like BIT2 were true, then A would not violate B’s right to privacy when she X-ray’s B’s safe again. This seems very strange.

²⁷ As we shall see, Rasmus Uhrenfeldt and I rejected this view in *Too Much Info*. But, for reasons we shall see in *Reply to Munch and Lundgren*, I have since changed my mind.

The Hybrid View, on the other hand, maintains that there is a reduction of B's privacy at t_3 , because A *perceives* the content of A's safe. On the Hybrid View, perception is a sufficient condition for a privacy reduction. Thus, the Hybrid View squares well with the normative intuition that A violates B's right to privacy at both t_1 and t_3 .

In this section, I have briefly defended what Munch and I have called the Hybrid View of privacy. Much more needs to be said in order to fully defend this view, but for present purposes, it suffices to only briefly sketch out the view, in order to prepare the ground for the arguments that I lay out in the following papers. Here is how the Hybrid View has bearing on the papers to follow: In *Too Much Info*, Rasmus Uhrenfeldt and I make a (perhaps questionable) distinction between descriptive- and normative versions of the Control Theory and the Access Theory. The Hybrid View can be seen as one instance of a descriptive Access Theory. In *Reply to Munch and Lundgren*, Rasmus Uhrenfeldt and I argue that it is possible to commit to the Control Theory when it comes to the *right* to privacy, while committing to the Access Theory when it comes to the *concept* of privacy. What this means is that it is possible to commit to, for instance, something like our Negative Control account of the *right* to privacy, while also committing to something like the Hybrid View of the *concept* of privacy. In *An Indirect Argument*, I indirectly defend the Access Theory of privacy. My defense of the Access Theory is not in conflict with the Hybrid View. Rather, the Hybrid View is an instance of the broader Access Theory. Let me briefly explain what

I mean by this. The version of the Access Theory that I defend in *An Indirect Argument* is the following:

The Access Theory

An individual A has informational privacy relative to another individual B and to a personal fact f about A if and only if B does not actually access f .

This version of the Access Theory is silent on what it means that B accesses f . The Hybrid View explicates exactly what it means that B accesses f , namely that B either forms a warranted belief in f , or that B perceives f .

This marks the conclusion of the chapter. Hopefully, all the previous sections in conjunction have prepared the ground for the papers to follow. The papers related to this chapter are: *Too Much Info*, *Reply to Munch and Lundgren*, and *An Indirect Argument*. The papers can be found at the end of the thesis, but here follows a brief summary of each of these papers.

3.5. Paper Summary: Too Much Info: Data Surveillance and Reasons to Favor the Control Account of the Right to Privacy

In *Too Much Info*, Rasmus Uhrenfeldt and I argue that there is at least a pro tanto reason to favor the Control Theory of the right to privacy over the rival Access Theory. The argumentative strategy in the paper is relatively simple, but also innovative: We first describe the Control Theory and the Access Theory in a way that seems uncontroversial, and in a way that their respective proponents would likely accept. Comparing the two theories at this stage does not reveal that any of them are clearly preferable to the other. Then we present each of the theories with three objections, and suggest how to improve each of the theories in order to escape the objections. When we arrive at the final versions of the theories – when we have strengthened the theories significantly – we introduce a test case to see which of the theories handles the test case best. The test case involves a paradigmatic example of a violation of the right to privacy. Thus, any plausible theory of the right to privacy must be compatible with the verdict that there is a violation of the right to privacy in the test case. As we argue, it turns out that the Control Theory handles the test case much more straightforwardly than the Access Theory does, and thus we have at least a pro tanto reason to favor the Control Theory over the Access Theory.

3.6. Paper Summary: Privacy Rights, and Why Negative Control is Not a Dead End: A Reply to Munch and Lundgren

Lauritz Munch and Björn Lundgren have published, independently of each other, reply papers to *Too Much Info* (also in Res Publica). In this paper, we respond to the main objections raised by Munch and Lundgren. Interestingly, they give almost identical counterexamples of our definition of Negative Control. We concede that the counterexamples are in fact genuine counterexamples, and we therefore give a new definition of Negative Control that is immune to the counterexamples. The definition we give draws heavily on the (non-normative version of the) definition of Negative Control that I had developed in the meantime in *An Indirect Argument*. With a few adjustments, the definition is immune to the objections raised by Munch and Lundgren.

In addition to the counterexample to our definition of Negative Control, Lundgren objects that we do not recognize that privacy is the object of the right to privacy. Lundgren claims that when we define the right to privacy in terms of control, as we do in *Too Much Info*, we are also committed to defining the concept of privacy in terms of control. This is a problem for us, according to Lundgren, because in *Too Much Info*, we do not consider the objections against control-based definitions of the concept of privacy. As a response to Lundgren, we explain how we think the concept

of privacy relates to the right to privacy. We concede that privacy is indeed the object of the right to privacy, but even so, we need not define the concept of privacy in terms of control just because we define the right to privacy in terms of control. We argue that when we say that ‘control’ should be defined as Negative Control, then this only means that the right to privacy is a certain type of control-right, similar to the control-rights that are conventionally taken to be a part of the bundle of rights that make up property rights.

3.7. Paper Summary: An Indirect Argument for the Access Theory of Privacy

In this paper, I develop the idea of Negative Control further, and I try to show what happens if we adopt a non-normative version of Negative Control. I show that if control theorists define ‘control’ as Negative Control, then all the standard objections to the Control Theory lose their bite. I go through a range of counterexamples to the Control Theory from the literature, and I show that all of them assume a definition that is either Positive Control or Republican Control. None of them cut any ice against the Control Theory, if control is defined as Negative Control.

However, it turns out that at the descriptive level, defining control as Negative Control collapses the Control Theory into the Access Theory. That is, the definition of Negative Control needed in order to resist all the objections to the Control Theory sneaks in notions of access in a way that implies that access is a necessary condition for a violation of the right to privacy.

I then discuss a recent version of the Control Theory developed by Leonhard Menges. He calls it the ‘source control account’. Menges’ account does *not* collapse into the Access Theory, so adopting his account seems to be a *prima facie* plausible option for the control theorist to avoid my collapse-objection. However, I show that Menges’ source control account should be rejected for other reasons. This leaves the

control theorists with my Negative Control account, which collapses into the Access Theory. Unless the control theorists can come up with a better theory that avoids all the classic objections to the Control Theory, the theory should be rejected in favor of the Access Theory.

CHAPTER 4. Data Ownership

It has often been suggested that one effective way to protect people’s informational privacy is to grant them property rights over personal data that pertains to them (Laudon, 1996); (Samuelson, 2000); (Thouvenin, Weber, & Früh, 2017); (Ritter & Mayer, 2018). If individuals legally own data about themselves, then they have a strong protection of the data, and legal measures can be taken if others violate their property rights. Among other suggested benefits, granting legal property rights over personal information gives individuals additional means to protect their privacy. The idea of data ownership seems *prima facie* appealing, and it has gained quite a lot of traction during the last few years. The idea also makes for a great slogan for companies that want to convince customers that it is safe to share personal information: “Data about you is yours!” Similarly, during a recent congressional hearing, Mark Zuckerberg claimed repeatedly that Facebook users “own all of their own content”.²⁸

So far, however, the consensus in legal theory seems to be that personal data cannot be owned (Hummel, Braun, & Dabrock, 2020). Not even the GDPR grants

²⁸ See <https://www.brookings.edu/blog/techtank/2019/06/26/why-data-ownership-is-the-wrong-approach-to-protecting-privacy/> (Accessed May 11, 2021).

property rights to individuals over personal information about themselves (Ibid). There has been a relatively large amount of criticism of the very idea of data ownership (Cohen, 2017); (Stepanov, 2020) (Purtova, 2009); (Schwartz, 2004); (Macnish & Gauttier, 2020); (Cofone, 2020). Nevertheless, the idea of data ownership is widely discussed in the political arena, and several attempts have been made to pass legislation that grants individuals property rights over personal information. For instance, Senator John Kennedy has introduced a bill called the ‘Own Your Own Data Act of 2019’.²⁹

The idea of data ownership has also spawned a series of new start-up companies, whose main purpose is to facilitate a market in personal information, in which the data subjects can participate directly. Companies like CitizenMe, and Datawallet enable individuals to sell personal data on a free market (Macnish & Gauttier, 2020: 43). The idea is that individuals can block the flow of data from themselves and to companies like Facebook and Google using specific types of web browsers, and instead let the data flow to companies like Datawallet. The company can either sell the data directly to other companies and give the data subjects a percentage of the profit, or they can pay the data subjects up front before they sell the

²⁹ See <https://www.govtrack.us/congress/bills/116/s806>. (Accessed May 11, 2021).

data to others. These data markets can take the form of monetary markets, barter markets, of combinations of the two.

The general idea of a market in personal information seems to assume that individuals own information about themselves. In order to sell something, you either have to own it yourself in the first place, or the owner must have agreed to let you sell it on her behalf (Macnish & Gauttier, 2020: 42). *Prima facie*, it does not make sense to talk of buying and selling personal information, if personal information is either not the sort of thing that can be owned at all, or if the individuals who are thought to be selling the information do not own it in the first place (or if they do not have the permission from the real owner to sell it).

I have recently defended the idea of markets in personal information in a paper that is not included in this thesis (Mainz, Forthcoming). In that paper, I argue that such markets will generally benefit the participants in it, and that the problems that may arise because of the market can be alleviated through proper regulation. However, I think it is possible to defend such markets without committing to the idea that individuals literally own personal data about them. Just like it is meaningful to talk about a ‘market in votes’ without committing to the view that voters literally have property rights in their votes, it is also meaningful to talk about a market in personal data without committing to the view that individuals literally have property rights in personal data (Freiman, 2014: 760). One way to think about what goes on in a vote

market is that individual A pays individual B to perform a specific action, namely to vote in a particular way (Ibid.). Similarly, we can say that what goes on in a market in personal information is that individual A pays B to perform a specific action, namely to give A access to particular pieces of information about B. I think the arguments in favor of allowing markets in personal information are overwhelming, but luckily, this view does not force me to accept the view that individuals should be granted property rights in personal information about them.

So far, we have been talking about property rights in personal information as something that we can ‘grant’ people by legal means. But perhaps people have *natural* property rights in personal information even if no current positive law recognizes this. Perhaps some of the classic philosophical theories of property rights imply that individuals own personal data about themselves. Francis Cheneval has recently defended this view in a paper published in *Critical Review of International Social and Political Philosophy*. Cheneval argues that the conclusion that people have property rights in personal information about themselves can be reached through two distinct strategies. The first one is to reach the conclusion through a standard Lockean theory of property. The second one is to reach the conclusion through a Rawlsian theory of distributive justice. In *Reply to Cheneval*, I try to explain why the first strategy fails, while I remain agnostic about the second strategy. In particular, I explain how Cheneval’s Lockean theory of property in personal information runs into a dilemma.

With these brief preliminary remarks, let us now turn to a summary of *Reply to Cheneval*, where I explain how the argument goes in more detail.

4.1. Paper Summary: But Anyone Can Mix Their Labor: A Reply to Cheneval

In this paper, I show how Francis Cheneval's Lockean argument for data ownership faces a dilemma. Cheneval's basic idea is that Locke's theory of property implies that people own personal data about themselves. I summarize Locke's argument as follows:

The Lockean Argument

Premise 1: If persons are the original owners of their respective personhoods, bodies and minds, then mixing their labor with something unowned generates property rights over the thing in question (provided that a certain proviso is satisfied).

Premise 2: Persons are the original owners of their respective personhoods, bodies and minds.

Conclusion: Mixing a person's labor with something unowned generates property rights over the thing in question (provided that a certain proviso is satisfied). (Mainz, 2021a: 278-279).

Cheneval seems to believe that applying either Premise 1 or Premise 2 of the Lockean Argument to personal data leads to the conclusion that people have property rights in personal data about themselves. ‘Applying Premise 1 to personal data’ means that individuals can come to own personal data, if they mix their labor with the data in certain ways. ‘Applying Premise 2 to personal data’ means that individuals own personal data about themselves, exactly because they own themselves. My main objection to Cheneval’s argument is that it leads to a dilemma, consisting of the following two options:

Option 1: Explain data ownership by applying Premise 1 in the Lockean Argument to personal data.

Option 2: Explain data ownership by applying Premise 2 in the Lockean Argument to personal data. (Mainz, 2021a: 282).

Option 1 has very counterintuitive implications. For instance, it implies that anyone can mix their labor with someone else’s personal data, and thus come to own personal data about others. If the data subject herself does not mix her labor with the data, then she might not even herself get any property rights over the data in question. This seems like a very strange implication. It seems strange that if other people process personal data about you before you do it, then they come to own the data. It seems especially strange given that Cheneval also believes that individuals own data about themselves

qua owning themselves. If Cheneval instead chooses Option 2, then he faces the following problem:

... then he cannot also defend data ownership through Premise 1, since an original owner of X does not lose any ownership in X just because someone else mixes labor with X. For this reason, if Cheneval chooses Option 2, then he loses his explanation for how people get partial ownership over personal data about other people, since his explanation consists in applying Premise 1 to personal data. (Mainz, 2021a: 280).

The idea is that if Cheneval chooses Option 2, then it follows that you already own the data about you. But in that case, others do not get *any* ownership of the data by processing it, because an original owner of something does not lose her property right just because someone mixes his labor with the thing in question. In the paper, I argue that the most plausible way out of the dilemma is to choose Option 2. However, simply applying the idea of self-ownership to personal data requires a lot of further theoretical work. Moreover, the solution is vulnerable to the objections against the idea of self-ownership as such. The overall conclusion of the paper is that Cheneval's Lockean defense of data ownership fails, unless he finds a plausible way of resolving the dilemma.

CHAPTER 5. Privacy & Inferences

Various types of data analytics make it possible to accurately infer personal information about individuals. When you give Facebook access to all sorts of information about you, including information about who your friends are, what you do in your spare time, what products you buy, which political demonstrations you attend, and so on, Facebook might try to train machine learning algorithms in order to infer ‘new’ information about you. As mentioned in the introduction, inferences of personal information are used for all sorts of purposes, and the inferences are made by private companies as well as state actors. There are at least two questions related to this practice that are relevant in the light of the previous chapters:

- i) Can a correct inference of an individual’s personal information *diminish* the individual’s privacy?
- ii) Can a correct inference of an individual’s personal information *violate* the individual’s right to privacy?

Let us start with i). Intuitively, it seems meaningful to say things like “Facebook diminishes my privacy when they infer my political preferences!” When Facebook infers your political preferences based on information that you voluntarily share, perhaps coupled with information that your friends and family share, Facebook comes

to ‘know’ something new about you that you can plausibly claim was a personal matter. You seem to have lost some privacy with regards to the information about your political preferences, and in relation to Facebook. How well does this square with what I have claimed so far in this thesis? It seems to square well with the version of the Access Theory that I have called the Hybrid View of privacy. Recall that on this view, there are two ways in which your privacy can be diminished: Someone else forms a warranted belief in a personal fact about you, or someone else perceives your personal matters. Now, in the specific case of Facebook inferring your political preferences, it may be a stretch to say that your privacy is diminished, if what really happens at Facebook is that no person or agent capable of forming warranted beliefs actually form such beliefs. It may be, for instance, that all of the inferential calculations are performed in a complex deep-learning algorithm the outputs of which no one ever accesses in an epistemic way, or ever perceives. Similarly, it may be that all of the inferential calculations are performed on completely anonymized data, so that even if a person or agent capable of forming warranted beliefs were to access ‘your’ data, they could not know that it was yours. I need not take a stance on any of these questions. What concerns us presently is the principled question of whether inferences as such –not necessarily made by a non-human entity like a Facebook algorithm – can diminish privacy.

It seems that there is no good reason to think that making an inference should *not* be one way to form a warranted belief about someone else. In *Inferences*, I use the following example: Jones owns a pickup truck, and Tim is the neighbor of Jones. Jones is proud of his car, and he frequently bores Tim with technical details about the car. Tim works as a data scientist. He wants to know what the correlations are between seemingly trivial data about electors, and their political preferences. He decides to find out whom Jones is likely to vote for in the upcoming election. He gets access to large amounts of data from publicly available databases, and trains a precise machine learning model on the data. To his surprise, Tim discovers that owning certain types of pickup trucks is a very strong predictor of voting Republican, and that owning certain types of sedans is a very strong predictor of voting Democrat. Based on all the technical details about the car that Tim has listened to in the driveway, he knows that Jones owns the exact type of pickup truck that correlates very strongly with voting Republican. It so happens that Jones in fact always votes Republican. Jones does not want Tim to know his political preferences, and he is not aware that it is possible to infer his political preferences based on information about which car he drives. Tim now asks the computer to calculate the likelihood of Jones voting Republican. Based on the correlations in the dataset, and the fact that Jones owns a specific type of pickup truck, the computer runs something like the following inference:

($\alpha 1$) Jones owns a pickup truck of type X.

(β 1) If one owns a pickup truck of type X, then one is very likely to vote Republican.

(γ 1) Jones is very likely to vote Republican.

It seems that Tim now has a warranted belief about Jones' political preference. If the Hybrid View of privacy is correct, then it follows straightforwardly that Jones' privacy has been diminished with regards to his political preference, and in relation to Tim. In general, it is difficult to see what should license the view that there is a relevant difference between Tim forming the belief about Jones' political preference in the way described above, as compared to forming the same belief based on, say, testimonial evidence from Jones himself (on the assumption that the probative values are more or less identical in the two cases).³⁰

Based on these brief remarks, it seems at least intelligible that inferences can diminish individuals' privacy. It thus seems that the answer to question i) is "yes". What about question ii)? Can inferences also *violate* individuals' privacy rights? This is the main question that I attempt to answer in *Inferences*. In the paper, I defend the

³⁰ Lauritz Munch has recently made a similar point, namely that it does not make a morally difference either (Munch, 2021).

view that an inference does *not* violate the right to privacy, if the pieces of information that the inference is based on are themselves obtained legitimate. I illustrate this by use of the example of Jones and Tim above. It seems that Tim makes the inference γ_1 based on information that he has obtained legitimate, namely α_1 and β_1 . But if all the steps that lead to Tim having the inferred information about Jones' political preference are themselves *legitimate*, how can it then be *illegitimate* for Tim to obtain the inferred information by making the inference?

However, in the paper I remain open to the view that *some* inferences can violate the right to privacy. It may be, for instance, that certain inferences that are based on information that are *not* obtained legitimate are themselves illegitimate. This also explains why it was initially important to answer question i). In an earlier section, I explained why privacy diminishments are necessary for privacy violations. So, given that I want to be open to the view that *some* inferences violate the right to privacy, I also need to be open to the view that at least *some* inferences can diminish privacy.

Much of the literature concerned with the question of whether inferences can violate privacy rights focus on examples involving harmful diminishments of privacy by inferring personal information about individuals. Indeed many such inferences can be harmful even if they do not violate anyone's right to privacy. However, I think it is an underappreciated fact that sometimes it can be harmful if an individual's privacy is *not* diminished through certain inferences. The cases I have in mind here are those

where decisional algorithms - like the ones discussed in the introduction – are used to decide whether people get a loan in the bank, get released on parole, or get the job they applied for. Due to mechanisms like privacy dependencies (also described in the introduction), it can be very difficult for minorities to break the spell of others' past behavior. For example, it can be very difficult for a black defendant to be released on parole because the data on *other* black defendants show that they have tended to recidivate when they were released. If you are a black defendant that as a matter of fact will not recidivate if released on parole, then you would want the justice system to infer *your* particular likelihood of recidivating, instead of basing the decision solely on the inferences of other black defendants' likelihood of recidivating.

5.1. Paper Summary: Inferences and the Right to Privacy

In this paper, I defend what I call the ‘Inference Principle’. This principle holds that if an agent obtains some information legitimately, then it is legitimate for the agent to make any inference based on the information. If, for example, Facebook were to obtain some trivial information about an individual legitimately, then it is legitimate for Facebook to make any inference based on the information, no matter how personal the inferred information is, and no matter if the individual has consented to Facebook making the inference. I try to show how the Inference Principle resembles Robert Nozick’s famous entitlement theory. The idea is that inferring personal information cannot be morally illegitimate if the steps that lead to someone having a piece of inferred information about someone else start from a legitimate baseline, and if all the steps are themselves legitimate. Nozick’s entitlement theory is widely considered controversial, but I try to show that even if the most common objections against it are correct, they do not have force against the Inference Principle. The general strategy of the paper is to build a positive defense of the Inference Principle, and then show that even if the positive defense is not completely compelling, the Inference Principle is at least more plausible than the views that contradict it.

I consider two objections to the Inference Principle. The first one comes from Benedict Rumbold, and James Wilson. Rumbold and Wilson have recently published a paper in which they defend a view that clashes with the Inference Principle. They

claim that waiving the right to privacy over some information (for instance by voluntarily making it public), does not imply waiving the right to privacy over any information that is inferred from it. The individual only waives her right to privacy over the inferred information, if she explicitly waives her right to privacy over it. I try to show why Rumbold and Wilson's view have very counterintuitive implications, and that it should ultimately be rejected.

Finally, I consider the objection that some inferences of personal information are other-regarding in a certain sense, and that the Inference Principle therefore proves too much. When privacy dependencies (as described in an earlier chapter) emerge, it is possible to infer information about individuals who did not contribute to the information on which the inference is based. So, when an inference about you is made, and you did not contribute to the information on which it is based, then it cannot be the case that you waived your right to privacy over the inferred information by waiving it over the information on which the inference is based. I try to explain why the objection misfires, and I try to show that even if the objection is correct, the Inference Principle can handle it in many real-world cases where online information is obtained in illegitimate ways.

5.2. Predicting Voter Behavior

As we saw in an earlier chapter, data analytics is used to infer many different types of personal information. In *Inferences*, argue that these inferences do *not* violate individuals' privacy rights, if the information on which the inferences are based were themselves obtained legitimately. In this section, we take a closer look at one particular type of personal information that can be inferred using state of the art data analytics: voter behavior. Some of the things described in this brief section will only become clear *after* reading *How to Buy an Election*. If helpful, I invite the reader to read the paper first, or at least the summary of it, and then return to this section.

Now, even if it does not violate privacy rights when someone infers the voting behavior of individual voters, it does not necessarily give a moral carte blanche to *use* the inferences in whatever way they feel like. The list of potentially problematic ways of using inferred information about individual voter behavior is long, but we shall focus on one particular way that makes it possible to legally buy an election in the US. Before turning to the paper *How to Buy an Election*, where Rasmus Uhrenfeldt, Jørn Sønderholm and I describe how data analytics can be used to buy an election, I shall first describe some of the empirical findings related to possibility of inferring the voting behavior of individuals.

The first thing to note about the literature on predicting voting behavior at an individual level is that the results are mixed. Some studies find that it is possible to predict individual voter behavior accurately using only very few trivial data points. Based entirely on Facebook ‘likes’, one study found that it was possible to accurately predict voter behavior, even in multi-party systems like the one in Denmark (Kristensen et al., 2017). Another study also found significant correlations between ordinary Facebook-users’ data, and their voting behavior in the 2016 US presidential election (Idan & Feigenbaum, 2019). However, a study conducted on German federal election data found no significant correlations between digital trace data and voting behavior (Bach et al., 2019).

There are several possible explanations for why the results are mixed. The first explanation is that the studies are based on very different datasets, containing different types of data and data of different quality. The second explanation is that the studies are carried out on very different populations, with very different political environments. In relatively polarized societies like the US with a de facto two-party system in place, it is probably not surprising that there are often robust correlations between trivial demographic information and voting behavior. In more politically diverse societies like Denmark, with a multi-party system in place it will *ceteris paribus* be more difficult to predict individual voter behavior. Furthermore, some of the studies are based on digital data collected *after* the implementation of the GDPR

in the EU. Article 17 of the GDPR gives data subjects the ‘right to be forgotten’, which means that the data subject can demand from a data controller (such as the company that collects the data) that data concerning the data subject is erased.³¹ In general, data protection regulations such as the GDPR may have certain chilling effects on the collection of personal data generated online (Sanders, 2019). With less personal data collected by companies, more anonymizations of the data, and more data subjects exercising their right to be forgotten, it may become more difficult to find strong correlations between digital data and voting behavior.

Despite the mixed empirical findings related to the possibility of accurately predicting voting behavior at an individual level, it is important to stress that the method for using data analytics to buy an election described in *How to Buy an Election* does not rest on the premise that the empirical findings are uniform. Neither does it rest on the premise that it is possible to consistently and accurately predict individual voter behavior at an individual level for all or nearly all voters. The reason for this is two-fold.

³¹ See <https://gdpr.eu/article-17-right-to-be-forgotten/> (Accessed May 11, 2021).

First, it *is* possible to accurately predict voting behavior at an individual level, at least for *some* groups of voters. To illustrate this point, consider a particular ‘voting prediction tool’ developed by The Economist. On their website, you can plot in 14 data points about an individual voter. The data points describe the voter’s skin-color, religion, income, marital status, parental status etc. Based on these data points, the tool will then tell you the likelihood of the voter in question voting for the Republicans, or the Democrats, respectively.³² Some combinations of data points will generate probabilities around .50, while other combinations will generate probabilities around .97. For instance, according to The Economist, there is .50 probability that an American voter that fits the following description will vote Democrat: a white straight woman who is a Catholic, is unmarried, has no children, has no college education, is between 45 and 64, earns \$30-64K a year, lives in the suburbs in the West, and does not speak Spanish. In contrast, there is .98 probability that an American voter that fits the following description will vote Democrat: a black straight woman who is an atheist, is married, has no children, has no college education, is 65 or older, earns less than \$30K a year, lives in a city in the West, and speaks Spanish. If the tool developed

³² See <https://www.economist.com/graphic-detail/2018/11/03/how-to-forecast-americans-vote> [Accessed May 11, 2021].

by The Economist is somewhat accurate, then it is indeed possible to predict the voting behavior of at least *some* voters. Anyone who is interested in buying an election in the way we describe in *How to Buy an Election* can simply focus on the demographic groups whose voting behavior is easier to predict.

Second, the method for buying an election that we describe in *How to Buy an Election* works even if the predictions are not very accurate. Suppose that individual K prefers candidate R over candidate D. K therefore decides to use the method for buying an election that we describe in *How to Buy an Election*. K hires a data analytics company to predict which voters are likely to vote for D. The most important thing for K is that she pays *enough* D-electors to abstain. In other words, K need not be able to know with absolute certainty that every given elector that she pays to abstain is a D-elector, in order for her to succeed with her efforts to buy the election. As long as the predictions pick out D-electors more often than it picks out R-electors, K's efforts can have a real effect on the election result. Of course, the higher the false positive rate is the more money it will cost K to win the election, because more actual D-electors will need to be paid to abstain in order to compensate for the false positives. With these brief remarks, let us now turn to the summary of *How to Buy an Election*.

5.3. Paper Summary: Big Data Analytics and How to Buy an Election

In this paper, Jørn Sønderholm, Rasmus Uhrenfeldt and I, show how it is possible to legally buy an election. The method we describe for buying an election involves the use of data analytics to predict how individual electors will vote in an upcoming election. Both the Democrats and the Republicans have access to huge amounts of data about individual voters. Based on the data, they can accurately predict how individual voters will vote in an upcoming election. Someone interested in buying an election can then approach the individual voters who are likely to vote for the opponent, with the offer of signing an employment contract. An implication of signing this contract and complying with it is that it becomes impossible for the voter in question to vote on Election Day. For instance, the contract could legally demand that the individual who signs it must be engaged in picking up trash in a specific local park that is located outside of the individual's home county, in exchange for being paid a certain amount of money.

In the US, there are publicly available voter registration lists in place, where one can see if a specific elector has voted (not *how* the elector voted). With access to these voter registration lists, it becomes possible to verify ex post if the individual who signs the contract breached the contract and voted anyway.

We go through a range of different policy responses that may block the possibility of buying an election in the way that we describe. For instance, if it were much easier to vote early, then fewer voters might be prevented from voting on Election Day by signing and complying with the contract. Similarly, if all employers had a right to vote during working hours, then signing and complying with the contract would not prevent the elector from voting on Election Day. Finally, if the voter registration lists were deleted, or at least made inaccessible, then it would be much more difficult to buy an election in the way we describe, because paying the elector hoping that she will comply with the contract and abstain from voting becomes nothing but a gamble. We go through these and other policy responses in turn, and discuss their respective likelihoods of effectively blocking the method for buying an election in the way we describe.

References

- Allen, A. (1988). *Uneasy access: Privacy for women in a free society*. Rowman & Littlefield Publishers.
- Allen, A. (2003). *Why privacy isn't everything: Feminist reflections of personal accountability*. Rowman & Littlefield Publishers.
- Altman, I. (1976). Privacy: A conceptual analysis. *Environment and Behavior*, 8(1), 141-141.
- Bach, R. L., Kern, C., Amaya, A., Keusch, F., Kreuter, F., Hecht, J., & Heinemann, J. (2019). Predicting voting behavior using digital trace data. *Social Science Computer Review*, Online first.
- Barnes, S. B. (2006). A privacy paradox: Social networking in the united states. *First Monday*, 9(11)
- Barocas, S., & Levy, K. (2020). Privacy dependencies. *Washington Law Review*, 95(2), 555-615.
- Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671-732.

- Barth, S., & Jong, Menno D. T. de. (2017). The privacy paradox – investigating discrepancies between expressed privacy concerns and actual online behavior – A systematic literature review. *Telematics and Informatics*, 34(7), 1038-1058.
- Beardsley, E. (1971). Privacy: Autonomy and selective disclosure. In R. Pennock, & J. Chapman (Eds.), *Privacy & personality* (pp. 56-70) Routledge.
- Benzanson, R. (1991). Privacy, personality, and social norms. *Case W. Res. L. Rev.*, 41(3), 681-687.
- Binns, R. (2018). Fairness in machine learning: Lessons from political philosophy. Paper presented at the Conference on Fairness, Accountability and Transparency, PMLR, 149-159.
- Blaauw, M. (2013). The epistemic account of privacy. *Episteme*, 10(2), 167-177.
- Blacklaws, C. (2018). Algorithms: Transparency and accountability. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2128) doi:10.1098/rsta.2017.0351
- Bok, S. (1989). *Secrets: On the ethics of concealment and revelation*. New York: Vintage Books.

Brendel, E. (2004). Intuition pumps and the proper use of thought experiments.

Dialectica, 58(1), 89-108.

Bunzl, M. (1996). The logic of thought experiments. *Synthese*, 106(2), 227-240.

Calo, M. R. (2011). The boundaries of privacy harm. *Indiana Law Journal*, 86(3),

1131-1162.

Castets-Renard, C. (2019). Accountability of algorithms in the GDPR and beyond: A

European legal framework on automated decision-making. *Fordham Intellectual*

Property, Media and Entertainment Law Journal, 30(1)

Castro, C. (2020). What's wrong with machine bias. *Ergo*, 6(15), 405-426.

Cofone, I. (2020). Beyond data ownership. *Ssrn*, doi:10.2139/ssrn.3564480

Cohen, J. (2017). Law for the platform economy. *University of California, Davis, Law*

Review, 51(133), 133-204.

Daniels, N. (1979). Wide reflective equilibrium and theory acceptance in ethics.

Journal of Philosophy, 76(5), 256–282.

Davis, F. (1959). What do we mean by right to privacy. *South Dakota Law Review*, 4

DeCew, J. (1997). *In pursuit of privacy: Law, ethics, and the rise of technology* Ithaca:

Cornell University Press.

Elster, J. (2011). How outlandish can imaginary cases be? *Journal of Applied Philosophy*, 28(3), 241-258.

Fairfield, J., & Engel, C. (2015). Privacy as a public good. *Duke Law Journal*, 65(3), 385-457.

Fallis, D. (2013). Privacy and lack of knowledge. *Episteme*, 10(2), 153-166.

Feinberg, J. (1986). *The moral limits of the criminal law volume 3: Harm to self*. Oxford University Press.

Freiman, C. (2014). Vote markets. *Australasian Journal of Philosophy*, 92(4), 759–774.

Fried, C. (1968). Privacy. *Yale Law Journal*, 77(3), 475-493.

Fried, C. (1984). Privacy: A moral analysis. In F. Schoeman (Ed.), *Philosophical dimensions of privacy: An anthology* (pp. 203-222) Cambridge University Press.

Gavison, R. (1980). Privacy and the limits of law. *The Yale Law Journal*, 89(3), 421-471.

Goldberg, I., Hill, A., & Shostack, A. (2001). Trust, ethics and privacy. *Boston University Law Review*, 81(2), 407-422.

Goldman, A. (1999). *Knowledge in a social world*. Oxford University Press.

Hargittai, E., Marwick, A. (2016). “What can I really do?” explaining the privacy paradox with online apathy. *International Journal of Communication*, 10

Hargittai, E., & Litt, E. (2013). New strategies for employment? internet skills and online privacy practices during people's job search. *IEEE Security Privacy*, 11(3), 38-45. doi:10.1109/MSP.2013.64

Hedden, B. (2021). On statistical criteria of algorithmic fairness. *Philosophy and Public Affairs*, 49(2), 209–231.

Hellman, D. (2020). Measuring algorithmic fairness. *Virginia Law Review*, 106(4), 811-866.

Henkin, L. (1974). Privacy and autonomy. *Columbia Law Review*, 74(8), 1410-1433.

Hill, K. (2012). How target figured out a teen girl was pregnant before her father did.

Retrieved from <https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/>

- Huemer, M. (2005). *Ethical intuitionism*. Palgrave Macmillan UK.
- Hummel, P., Braun, M., & Dabrock, P. (2020). Own data? ethical reflections on data ownership. *Philosophy & Technology*, Retrieved from <https://doi.org/10.1007/s13347-020-00404-9>
- Idan, L., & Feigenbaum, J. (2019). Show me your friends, and I will tell you whom you vote for: Predicting voting behavior in social networks. Paper presented at the *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, Vancouver, British Columbia, Canada. 816–824.
- Inness, J. C. (1992). *Privacy, intimacy, and isolation*. Oxford University Press.
- Kahneman, D. (2013). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Kappel, K. (2013). Epistemological dimensions of informational privacy. *Episteme*, 10(2), 179-192.
- Kim, T. W., & Routledge, B. (2020). Why a right to an explanation of algorithmic decision-making should exist: A trust-based approach. Retrieved from <https://papers.ssrn.com/abstract=3716519>

- Knight, C. (2017). Reflective equilibrium. In A. Blau (Ed.), *Method in analytical political theory* (pp. 46-64) Cambridge University Press.
- Kristensen, J. B., Albrechtsen, T., Dahl-Nielsen, E., Jensen, M., Skovrind, M., & Bornakke, T. (2017). Parsimonious data: How a single facebook like predicts voting behavior in multiparty systems. *Plos One*, 12(9), e0184562.
- Kuyoro, A., Goga, P. N., Awodele, D. O., & Okolie, D. S. (2013). Optimal algorithm for predicting students academic performance. *International Journal of Computers & Technology*, 4(1), 63-75. doi:10.24297/ijct.v4i1b.3061
- Laudon, K. (1996). Markets and privacy. *Communications of the ACM*, 39(9), 92-104.
- Lippert-Rasmussen, K. (2008). Against self-ownership: There are no fact-insensitive ownership rights over one's body. *Philosophy and Public Affairs*, 36(1), 86–118.
- Lippke, R. L. (1989). Work, privacy, and autonomy. *Public Affairs Quarterly*, 3(2), 41-55.
- Lundgren, B. (2020). A dilemma for privacy as control. *Journal of Ethics*, 20, 165-175.
- Macnish, K. (2018). Government surveillance and why defining privacy matters in a post-snowden world. *Journal of Applied Philosophy*, 35(2), 417-432.

- Macnish, K., & Gauttier, S. (2020). A pre-occupation with possession: The (non-) ownership of personal data. In K. Macnish, & S. Gauttier (Eds.), *Big data and democracy* (pp. 42-54) Edinburgh University Press.
- Madden, M., & Rainie, L. (2015). Americans' attitudes about privacy, security and surveillance.
- Mainz, J. (2021a). But anyone can mix their labor: A reply to cheneval. *Critical Review of International Social and Political Philosophy*, 24(2), 276-285.
- Mainz, J. (2021b). An indirect argument for the access theory of privacy. *Res Publica*, Retrieved from <https://doi.org/10.1007/s11158-021-09521-4>
- Mainz, J. (Forthcoming). Are markets in personal information morally permissible? *Journal of Information Ethics*,
- Mainz, J., & Uhrenfeldt, R. (2021). Too much info: Data surveillance and reasons to favor the control account of the right to privacy. *Res Publica*, 27(2), 287-302.
- Margulis, S. T. (1977). Conceptions of privacy: Current status and next steps. *Journal of Social Issues*, 33(3), 5-21.
- Marmor, A. (2015). What is the right to privacy? *Philosophy & Public Affairs*, 43(1), 3-26. doi:10.1111/papa.12040

- Matheson, D. (2007). Unknowableness and informational privacy. *Journal of Philosophical Research*, 32, 251-267.
- McMahan, J. (2013). Moral intuition. In H. LaFolette, & I. Persson (Eds.), *The blackwell guide to ethical theory* (pp. 103-120) John Wiley & Sons, Ltd.
- Meijer, A., & Wessels, M. (2019). Predictive policing: Review of benefits and drawbacks. *International Journal of Public Administration*, 42(12), 1031-1039.
doi:10.1080/01900692.2019.1575664
- Menges, L. (Forthcoming). A defense of privacy as control. *The Journal of Ethics*,
- Mokrosinska, D. (2018). Privacy and autonomy: On some misconceptions concerning the political dimensions of privacy. *Law and Philosophy*, 37, 117-143.
- Moore, A. (2008). Defining privacy. *Journal of Social Philosophy*, 39(3), 411-428.
- Moore, A. (2010). *Privacy rights: Moral and legal foundations*. Pennsylvania State University Press.
- Munch, L. (2020). The right to privacy, control over self-presentation, and subsequent harm. *Journal of Applied Philosophy*, 37(1), 141-154.

- Munch, L. (2021). Privacy rights and ‘naked’ statistical evidence. *Philosophical Studies*, Retrieved from <https://doi.org/10.1007/s11098-021-01625-0>
- Munch, L., & Mainz, J. (2021). To believe or not to believe, that is not the (only) question: A hybrid view of privacy. *Unpublished*.
- Noorhannah, B., & Jayabalan, M. (2018). Risk prediction in life insurance industry using supervised learning algorithms. *Complex & Intelligent Systems*, 4, 145-154.
- Parent, W. (1983). Privacy, morality and the law. *Philosophy and Public Affairs*, 12(4), 269–288.
- Parker, R. (1974). A definition of privacy. *Rutgers Law Review*, 27, 275-296.
- Pessach, D., & Shmueli, E. (2020). Algorithmic fairness. *arXiv*, Retrieved from <https://arxiv.org/abs/2001.09784v1>
- Prosser, W. L. (1960). Privacy. *California Law Review*, 48(3), 383-423.
- Purtova, N. (2009). Property rights in personal data: Learning from the american discourse. *Computer Law & Security Review*, 25(6), 507-521.

Rachels, J. (1975). Why privacy is important. *Philosophy & Public Affairs*, 4(4), 323-333.

Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2020). Mitigating bias in algorithmic hiring: Evaluating claims and practices. Paper presented at the *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, Barcelona, Spain. 469–481.

Rawls, J. (1971). *A theory of justice*. Harvard University Press.

Raz, J. (1988). *Autonomy and pluralism*. Oxford University Press.

Reiman, J. H. (1995). Driving to the panopticon: A philosophical exploration of the risks to privacy posed by the highway technology of the future. *Privacy* (pp. 159-176)

Ritter, J., & Mayer, A. (2018). Regulating data as property: A new construct for moving forward. *Duke Law & Technology Review*, 16(1), 220-277.

Roessler, B. (2005). *The value of privacy*. Polity.

Samuelson, P. (2000). Privacy as intellectual property? *Stanf Law Rev*, 52(5), 1125-1173.

- Sanders, A. K. (2019). The GDPR one year later: Protecting privacy or preventing access to information essays. *Tulane Law Review*, 93(5), 1229-1254.
- Scanlon, T. (1975). Thomson on privacy. *Philosophy and Public Affairs*, 4(4), 315-322.
- Schwartz, P. M. (2004). Property, privacy, and personal data. *Harvard Law Review*, 117(7), 2056-2128.
- Singer, P. (2005). Ethics and intuitions. *The Journal of Ethics*, 9(3/4), 331-352.
- Sinnott-Armstrong, W., Young, L., & Cushman, F. (2010). Moral intuitions. In J. Doris (Ed.), *The moral psychology handbook* (pp. 246-272). New York, NY, US: Oxford University Press.
- Skopek, J. (2020). Untangling privacy: Losses versus violations. *Iowa Law Review*, 105(5), 2169-2231.
- Solove, D. J. (2008). *Understanding privacy*. Harvard University Press.
- Stepanov, I. (2020). Introducing a property right over data in the EU: The data producer's right – an evaluation. *International Review of Law, Computers & Technology*, 34(1), 65-86. doi:10.1080/13600869.2019.1631621

- Susser, D., Roessler, B., & Nissenbaum, H. (2019). Technology, autonomy, and manipulation. *Internet Policy Review*, 8(2), 1-22.
- Taddicken, M. (2014). The 'Privacy paradox' in the social web: The impact of privacy concerns, individual characteristics, and the perceived social relevance on different forms of self-disclosure. *Journal of Computer-Mediated Communication*, 19(2), 248-273.
- Taylor, J. S. (2002). Privacy and autonomy: A reappraisal. *Southern Journal of Philosophy*, 40(4), 587–604. doi:10.1111/j.2041-6962.2002.tb01918.x
- Thomson, J. J. (1975). The right to privacy. *Philosophy and Public Affairs*, 4(4), 295-314.
- Thouvenin, F., Weber, R. H., & Früh, A. (2017). *Data ownership: Taking stock and mapping the issues*. CRC Press.
- Turkson, R. E., Baagyere, E. Y., & Wenya, G. E. (Sep. 2016). A machine learning approach for predicting bank credit worthiness. Paper presented at the 1-7. doi:10.1109/ICAIPR.2016.7585216
- van den Haag, E. (1971). On privacy. *Privacy: Nomos XIII* (pp. 149-168)

Warren, S., & Brandeis, L. (1890). The right to privacy. *Harvard Law Review*, 4(5), 193-220.

Westin, A. F. (1970). *Privacy and freedom*. Bodley Head.



Too Much Info: Data Surveillance and Reasons to Favor the Control Account of the Right to Privacy

Jakob Thrane Mainz¹ · Rasmus Uhrenfeldt¹

Published online: 16 July 2020
© Springer Nature B.V. 2020

Abstract

In this paper, we argue that there is at least a *pro tanto* reason to favor the control account of the right to privacy over the access account of the right to privacy. This conclusion is of interest due to its relevance for contemporary discussions related to surveillance policies. We discuss several ways in which the two accounts of the right to privacy can be improved significantly by making minor adjustments to their respective definitions. We then test the improved versions of the two accounts on a test case, to see which account best explains the violation that occurs in the case. The test turns out in favor of the control account.

Keywords Privacy rights · Surveillance · Ethics of surveillance · Control account · Access account

Introduction

This paper is about the right to privacy. We offer a range of specific suggestions as to how the two most popular accounts of the right to privacy can be improved, by adjusting their respective definitions slightly. The first account is the Control Account (CA), and the second is the Access Account (AA).¹ We will call the proponents of these accounts ‘control theorists’, and ‘access theorists’ respectively. After

¹ The CA, broadly conceived, has been developed by Warren and Brandeis (1890), Westin (1970), Fried (1968), Moore (2003, 2010), Gross (1971), Parker (1974), Parent (1983), Allen (2003), Rössler (2005), Bezanson (1992), Goldberg et al. (2001), Altman (1976), Ryan and Calo (2010), Margulis (1977), Miller (1971), Scanlon (1975), Inness (1992), and many more. The AA, broadly conceived, has been developed by Thomson (1975), Gavison (1980), Bok (1989), Allen (1988), van den Haag (1971), Macnish (2018), and others. Note that some theorists have contributed to both.

✉ Jakob Thrane Mainz
jtm@hum.aau.dk

Rasmus Uhrenfeldt
ru@hum.aau.dk

¹ Aalborg University, Kroghstræde 3, 9220 Aalborg, Denmark

having improved the accounts, we test them on a thought example to see which account best explains the violation in the example. This reveals a *pro tanto* reason to favor the CA over the AA.

There are both *descriptive* and *normative* versions of both accounts. Descriptive accounts explain the necessary and/or sufficient conditions for *having* or *losing* privacy. The normative accounts explain the necessary and/or sufficient conditions for *violations* of the moral *right* to privacy (whatever that is) to occur.² A descriptive account is, as Adam Moore suggests, an account that describes *a state* or *condition* of privacy while normative accounts refer to moral obligations and rights (Moore 2008, pp. 212–213). Imagine that an individual invites strangers to observe her while she is at home. This individual is now in a lessened *state of privacy*, but since she herself invited the observers, her right to privacy has not been violated.

In this paper, we focus on the normative accounts, unless specified otherwise.³ According to the control theorists, control is a crucial feature of the right to privacy. If, and only if, I lose control over access to the relevant information,⁴ is my right to privacy violated. The access theorists, on the other hand, argue that a loss of control of the access to the information in question is not sufficient for a violation of the right to privacy to occur. They argue that the information in question must also in fact be accessed, in order for the right to privacy to be violated.

When we say that a person has a right to privacy, we do not subscribe to any particular theory of what it means to have a right to something. All our arguments are compatible with all of the most common theories of rights. For example, according to the interest theory of rights, the function of a person's right to privacy is that having such a right furthers her interests. According to the will theory of rights, on the other hand, the function of a person's right to privacy is to give that person control over the duties of other persons with regards to her privacy. Since nothing in our arguments hangs on which account of rights is the correct one, we will remain agnostic about this. However, we will assume—uncontroversially—that a right to privacy is a waivable, non-absolute right.

Why does it matter whether the control or the access account of the right to privacy is the correct one? As the access theorist Kevin Macnish has recently pointed out, it matters a great deal for our normative evaluations of many cases related to surveillance. For example, it matters for our evaluation of the case of the National Security Agency (NSA) collecting significant amounts of personal data about American citizens, and our evaluation of Edward Snowden's revelations of this practice (Macnish 2018, p. 2). It seems that if the CA is correct, millions of citizens' right to privacy is violated when the NSA collects data about them. This is so, because the citizens lose control over the access to information about them. If, on the other hand, the AA is correct, then it seems that citizens' right to privacy has not been violated

² We write 'moral right' to distinguish it from a legal right.

³ It is frustratingly difficult to determine which accounts are meant to be descriptive, which are meant to be normative, and which are both. Among the theorists we discuss in this paper, we count Adam Moore's account as a normative CA, and Judith Jarvis Thomson's and Kevin Macnish's accounts as normative AAs.

⁴ In this paper, we focus on *informational* privacy, although many have argued persuasively that privacy also concerns other things like spaces or bodies (See e.g. Moore 2010, pp. 25–26).

by this practice. The right to privacy has been violated on the AA only if persons at the NSA (or others) actually access the information (ibid.). So, this is not only an interesting theoretical discussion about definitions. It potentially has very important and wide-reaching implications for national security policy and surveillance policy.

The paper is structured as follows. In ‘[The Control Account of the Right to Privacy](#)’ section, we provide an initial definition of the CA. In ‘[The Access Account of the Right to Privacy](#)’ section, we provide an initial definition of the AA. In ‘[Improving the CA and the AA](#)’ section, we present and discuss several ways in which the definitions of the two accounts can be improved. We then provide a test case, to see which account, in its improved version, best explains the violation in the test case. Finally, in ‘[Concluding Remarks](#)’ section we make a few concluding remarks.

The Control Account of the Right to Privacy

Let us now consider a definition of the CA. There is no universal consensus among the control theorists about how exactly the CA should be defined. The key idea is, though, that a loss of control over access to the relevant information is necessary and sufficient for privacy violations to occur. The definition we will provide in this section is meant to capture what most control theorists would subscribe to. In the following section, we will then try to improve this initial definition on the control theorists’ behalf. The initial definition which seems to catch the crux of what most control theorists have in mind is this:

CA1: For any agent A to have her right to privacy violated, there is a necessary and sufficient condition that must be satisfied: Agent A has involuntarily lost control over unwanted access to personal information P about agent A.⁵

We do not suggest that all control theorists use the exact wording of CA1.⁶ But we do think that any control theorist needs to accept that losing control over access to personal information is a necessary and sufficient condition for a violation of the right to privacy to occur. Otherwise, such a theorist does not count as a (normative) control theorist.

Moore is an example of a recent and prominent control theorist. According to Moore, ‘*A right to privacy is a right to control access to and uses of—places, bodies, and personal information*’ (Moore 2010, p. 27). As this quote indicates, Moore thinks that privacy is not exclusively concerned with *informational* privacy. His definition also covers ‘locational privacy’ and ‘physical privacy’, and it not only covers

⁵ Some control theorists do not include the access-part. See Schoeman (1984, pp. 2–3).

⁶ Despite the fact that it is very difficult to determine which control theorists think of their respective accounts as normative accounts, we think it is fair to say that the CA1 can at least be distilled from the accounts of Allen (1999), Parker (1974), and Moore (2008), but probably many more. Allen, for example, writes: “‘privacy’ means personal data control or rights of data control; that the right of privacy is a right of personal data control; and that enhancing personal data control by individuals is the optimal end of privacy regulation’ (p. 875).

control rights, but also *use*-rights. Nonetheless, Moore seems to endorse the CA1 when it comes to informational privacy.

To illustrate that, for Moore, a loss of control over access is a *sufficient* condition for a violation of the right to privacy to occur, he provides two cases:

Zone Intrusion: Suppose you look in my safe with your X-ray device to see what it holds—there could be a stolen photo, a borrowed photo, or nothing....

Mere Zone Intrusion: Just like the first zone intrusion case although the person looking has no short-term memory and will forget any fact learned immediately.

In the case of zone intrusion a right to control access has been violated even though nothing except a bare fact has been seized. This is further illustrated by the example of mere zone intrusion. In the second case, nothing has been taken—no facts have been learned—all that has happened is that a zone or boundary has been unjustifiably crossed. (Moore 2003, p. 423)⁷

In Mere Zone Intrusion, Moore thinks that a violation of the right to privacy has occurred, because control over access to information has been lost.⁸ The loss of control over access is thus sufficient for the violation of the right to privacy to occur.

To illustrate that, for Moore, a loss of control over access is also a *necessary* condition, consider the following case:

The Loud Fight: Suppose that Fred and Ginger are having a fight - shouting at each other with the windows open so that anyone on the street can hear. (Moore 2003, p. 421)⁹

Moore thinks that no violation of the right to privacy has occurred in The Loud Fight:

In the loud fight case it would seem that Fred and Ginger have waived the right to privacy - they have via their actions allowed others who are in a public space to hear the fight. (Moore 2003, p. 421)

In The Loud Fight, information has been accessed by the people on the street, but no violation of the right to privacy has occurred, according to Moore. Fred and Ginger still have control over the people on the street's access to the information, because Fred and Ginger could simply choose to close the windows. Moore thus thinks that a loss of control of the access is a necessary condition for the violation of the right to privacy to occur. Similar quotes can be found in the works of other control theorists, but we will let Moore serve as a canonical example.

⁷ Moore borrows the Zone Intrusion case from Thomson (1975, p. 298).

⁸ One might argue that information has indeed been accessed, although the person forgets the information immediately. That might be so, but it seems that this is not what drives Moore's intuition that a violation has occurred. What drives his intuition seems to be that control over access has been lost.

⁹ Moore borrows The Loud Fight case from Thomson (1975, p. 296).

The Access Account of the Right to Privacy

Let us now turn to the AA. The key motivator for the access theorists seems to be that the CA is too broad, since it, counterintuitively, allows for violations of the right to privacy in cases where control has been lost, but no actual access to information has occurred (Thomson 1975, p. 305). The access theorists therefore add the extra necessary condition that the information in question must actually be accessed, in order for a violation of the right to privacy to occur. The definition we will provide in this section is meant to capture what most access theorists would subscribe to. In the following section, we will then try to improve this initial definition on the access theorists' behalf. The initial definition which seems to catch the crux of what most access theorists to have in mind is this:

AA1: For any agent A to have her right to privacy violated there are two each necessary and jointly sufficient conditions that must be satisfied: (1) Agent A has involuntarily lost control over unwanted access to personal information P about A, and (2) agent B (or someone else) actually accesses P.

Understood this way, the AA adds a necessary condition to the CA, namely the condition (2). We do not suggest that all access theorists use the exact wording of AA1. But we do think that any access theorist needs to accept that losing control over access to personal information—in conjunction with actual access to this information—are each necessary and jointly sufficient conditions for a violation of the right to privacy to occur.¹⁰ Otherwise, such a theorist does not count as an (normative) access theorist.

Kevin Macnish has put it this way: '*In contrast to the control account, the access account holds that information needs to be accessed for there to be an actual violation of privacy*' (Macnish 2018, p. 4). Macnish is an example of an access theorist who defends the view that access is a necessary condition for a violation of the right to privacy to occur.¹¹ In order to demonstrate this point, Macnish provides the following example:

... imagine that I leave my diary on a table in a coffee shop and return to that shop 30 min later to retrieve it. When I enter the shop I see a stranger with my diary on her table, a different table from the one at which I was sitting. I therefore know that she, or someone, has moved my diary, but have they read it? I have not been in control of my diary for half-an-hour, in which time anything might have happened to it. (Macnish 2018, p. 4)

¹⁰ This means that access and control accounts overlap in some cases. This is so because the AA *adds* a necessary condition to the CA.

¹¹ It is very difficult to determine which access theorists think of their respective accounts as normative accounts, but according to the access theorist Macnish the position that access is necessary for a violation of the right to privacy to occur is held by Allen (1988), Bok (1989), Gavison (1980), Gross (1971), Thomson (1975), and van den Haag (1971). The AA1 can at least be distilled from the accounts of these theorists.

In this diary case, there is definitely a loss of control over access to the information in the diary. But, according to Macnish, no violation of the right to privacy has occurred. In order for a violation to occur, someone must open the diary and read it. If no one does so, no violation has occurred:

Imagine that I have returned to the coffee shop after a 30 min interval to find my diary on the table. It is unopened. 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. I feel an enormous sense of relief, thank her and leave with my dignity intact. In this case, I do not think that my privacy has been lessened. When I see my diary in another's possession, I fear that my privacy has been violated, and indeed it might have been. However, as long as the diary is not actually opened and read no reduction in privacy has occurred. Note that this is true even though the diary was not under my control for 30 min. (Macnish 2018, pp. 4–5)

Note that Macnish writes that privacy has been neither 'lessened', 'violated', nor 'reduced' in this quote. We interpret this to mean that Macnish thinks that the diary example applies to *both* the descriptive *and* the normative AA. So, in relation to the normative AA, Macnish seems to think that, in addition to a loss of control over access, the information in question must be accessed in order for a violation of the right to privacy to occur. Similar quotes can be found in the works of other access theorists, but we will let Macnish serve as a canonical example.

It is important to stress that on the AA, there must be an *actual* access, and not just an *ability* to access, in order for there to be a violation of the right to privacy. Alan Rubel has suggested the following rough summarization of the descriptive version of AA: 'Privacy has to do with others' actual access, or *ability* to access, a person' (Rubel 2011, p. 296 [our emphasis]). In relation to the *normative* version of the AA, it seems that access must be interpreted solely as *actual* access, and not the *ability* to access. The reason is that the latter seems to be similar to a lack of *control* on the claimant's side, which will collapse the AA into something close to the CA. If Jones has the *ability* to access information about Smith, but chooses not to make use of it, then in some way, Smith does not have *control* over the access. Jones's ability to access is a sufficient condition for Smith not having control over the access. Conversely, Smith not having control over the access is a necessary condition for Jones having the ability to access the information. In order to distinguish their position sufficiently from the CA, the access theorists therefore need to include only *actual* access in their definition.¹²

¹² By 'actual access', the access theorists seem to mean something like 'actual *epistemic* access'. A person must have formed an epistemic relation to the information in question in order for actual access to obtain.

Improving the CA and the AA

Given that the only thing that distinguishes the two accounts is an extra necessary condition in the AA, it is not surprising that much of the criticism that applies to one of the accounts, also applies to the other. This will be evident throughout this section, when we address new issues, some of which apply to both accounts, and suggest ways to accommodate these issues by making adjustments to the definitions. The first issue we will discuss concerns the meaning of the word ‘control’.

Positive Control Versus Negative Control

The word ‘control’ seems to mean different things to different people in the privacy literature. The plausibility of the CA and the AA depends to a significant extent on which interpretation of control is at play. Let us introduce a distinction between ‘Positive Control’, ‘Negative Control’:¹³ and ‘Republican Control’:¹⁴

Positive Control: Agent A enjoys Positive Control over the access to relevant information P, if, and only if, A tries (or could try) to give agent B actual access to P, and succeeds.

Negative Control: Agent A enjoys Negative Control over access to relevant information P, if, and only if, A is capable of preventing agent B, who attempts to access, from accessing P.

Republican Control: Agent A enjoys Republican Control if, and only if, agent B does not have the ability to get access to relevant information P about A.¹⁵

Only the distinction between Positive Control and Negative Control is of relevance for this section. Later, we will explain how the distinction between Negative Control and Republican Control offers an effective rejoinder to Judith Jarvis Thomson’s famous objection against the Control Account, and against Macnish’s diary case introduced in the previous section.

Let us first make a point about the definition of Negative Control. It is tempting to think that the definition of Negative Control implies that any loss of Negative Control results in an access of information, since a loss of Negative Control always comes with an attempt to access. This would make it difficult to conceptually separate the CA from the AA. However, as we shall see in ‘A Test Case’ section, there are cases in which the lack of access is due to contingent circumstances, and in such

¹³ The distinction between Positive Control and Negative Control is inspired by Isaiah Berlin’s famous distinction between ‘positive liberty’ and ‘negative liberty’ (Berlin 1969, pp. 121–122). However, there is a crucial difference: negative liberty has a contrafactual definition, while Negative Control does not.

¹⁴ This is inspired by Philip Pettit’s idea of ‘republican freedom’. See Pettit (1999).

¹⁵ We are not the first ones to consider the combination of republicanism and privacy. See, for example, Newell (2018), Roberts (2014), van der Sloot (2018), and Hoye and Monaghan (2015). However, all of these authors write about how privacy is important for retaining republican freedom. Our idea is different. We interpret control in a republican manner in order to improve the control account, so that it can escape certain objections.

cases, Negative Control over certain information can be lost, while no access to that information occurs.

Let us now try to explain why the distinction between Positive Control and Negative Control is important. Our claim is that *if* the control account should be taken seriously, it must explain all violations of the right to privacy in terms of Negative Control, and only Negative Control. A loss of Positive Control cannot plausibly violate the right to privacy. To see why, consider the following example:

Too Much Info #1

Suppose that Smith and Jones are co-workers. Smith likes to share personal information about his sex life. One day, as Smith is about to tell Jones something personal again, Jones simply puts his fingers in his ears before Smith starts talking. Smith finishes his story anyway.

If control is interpreted as Positive Control, Jones has violated Smith's right to privacy by putting his fingers in his ears, since Smith then loses Positive Control over the access to the information. But clearly, it would be absurd to maintain that a violation of the right to privacy has occurred in Too Much Info #1.¹⁶ Nonetheless, the interpretation of control as Positive Control can be found in the works of prominent privacy scholars, although they have not used the term 'Positive Control'. Take for example Jeffrey Reiman's use of the term in his influential critique of the control theorist Charles Fried: '*... in our culture one does not have control over who gets to observe one's performance of the excretory functions, since it is generally prohibited to execute them in public*' (Reiman 1995, p. 30).¹⁷ Here, it seems, Reiman's interprets control as Positive Control. Reiman's point is that if person A wants person B to have access to person A's performance of the excretory functions, but person A does not succeed, then person A lacks a relevant form of control. Contrast this form of control with the one in Moore's 'Zone Intrusion' and 'Mere Zone Intrusion' in '[The Control Account of the Right to Privacy](#)' section. In Moore's cases, control seems to be interpreted as Negative Control.

In order to avoid the strange implication of Too Much Info #1, and thus avoiding an accusation of the CA being too broad, the control theorist might want to specify that control should be interpreted as Negative Control, and only Negative Control. So, a revised definition of the CA could look like:

CA2: For any agent A to have her right to privacy violated, there is a necessary and sufficient condition that must be satisfied: Agent A has involuntarily lost *negative* control over unwanted access to personal information P about agent A.

On CA2, no violation occurs in Too Much Info #1, since no one has lost Negative Control. Note that Too Much Info #1 is not a problematic counterexample for

¹⁶ Joel Feinberg has argued that in a case like this, Smith has actually violated Jones's right to privacy by divulging private information unto Jones (Feinberg 1985, p. 23). As Feinberg would probably agree, this hinges on an interpretation of privacy, which conflates privacy with liberty or autonomy.

¹⁷ For another example, see Farber (1993, p. 515).

the access theorist, since Jones does not get access. However, we can make a slight alteration of Too Much Info #1, so that Jones does in fact get access. Call this example Too Much Info #2:

Too Much Info #2

The same as Too Much Info #1, but this time Peter has been standing in the same room as Smith and Jones without anyone noticing. After Smith has left the room, Peter tells Jones what Smith was trying to tell.

If control is interpreted as Positive Control, there is a violation in Too Much Info #2 on the AA1. This is so, because Smith does not have Positive Control over whether Jones has access, and Jones does in fact access the information. But it seems very implausible that Jones has violated Smith's right to privacy in Too Much Info #2. It seems more plausible that *Peter* violates at least Smith's right to privacy, due to Peter's eavesdropping. This violation can be explained as a loss of Negative Control on Smith's part. In order to rule out Positive Control, we suggest the same adjustment to the definition of the AA, as we did to the definition of the CA:

AA2: For any agent A to have her right to privacy violated there are two each necessary and jointly sufficient conditions that must be satisfied: (1) Agent A has involuntarily lost *Negative Control* over unwanted access to personal information P about A, and (2) agent B has access to P.

Note that A does not need to lose *all* of her Negative Control over P in order for her right to privacy to be a violation. A can have full Negative Control with regards to some agents, while having lost Negative control with regards to others. To see this point, consider Futuria.

Futuria

In Futuria each person at the age of 20 is forced by law to let one of 50 private companies have access to certain very personal information. Sarah has just turned 20 and therefore needs to choose which of these companies she wants to give her information to. She actively dislikes 48 of the companies and therefore uses her Negative Control to withhold her information from these companies. She is agnostic about giving her information to the remaining two companies, so she chooses one at random.

In Futuria, it seems that Sarah enjoys a substantial degree of control, but her right to privacy is still violated. The reason is that Sarah does not have control with regards to *all* of the companies. We cannot point to any of these companies and say 'Sarah was coerced to give information to this particular company'. However, what matters is if Sarah is in control over whether *any* agent has access to information about her.¹⁸

¹⁸ Note that this implies that many modern democratic states are constantly engaged in infringing on privacy rights, when relevant state authorities gain access to personal finances, medical records, etc. People may have differing intuitions in this case. Our intuition is that such states do in fact infringe on people's right to privacy, but that doing so can be justifiable on weightier non-privacy related grounds.

Wanted Versus Unwanted Access

Both control theorists and access theorists have claimed that privacy is concerned with *unwanted* access. In fact, no one in the literature seems to dispute this. Consider this quote from the control theorist Beate Rössler:

Something counts as private, if one can oneself control the access to this ‘something’. Conversely, the protection of privacy means protection against unwanted access by other people. (Rössler 2005, p. 8)

Or, this quote from the access theorist Sissela Bok: ‘The condition of being protected from unwanted access by others—either physical access, personal information, or attention’ (Bok 1989, pp. 10–11).

The notion of ‘unwanted’ can be spelled out in at least two different ways: (1) Either as a description of some actual or possible psychological state, such as the *absence* of a desire, ambition, unconscious or conscious wish (or an *active* disfavoring of this psychological state), or (2) as a more abstract normative concept, which is supposed to do some normative work on its own.

Let us first explain why it is not conceivable to understand unwanted as (2). On (2), unwanted is supposed to do some normative work on its own, and presumably refer to the importance of being able to exclude someone from having access. But on that interpretation, it is a bit unclear what normative work it does that is relevantly different from what the control theorists mean by (negative) ‘control’: If Smith has full control over the access, and Jones has access, it must at least be the case that Jones’s access is not unwanted by Smith. So, we assume that unwanted should be understood as (1) or something close to it.

If we understand unwanted as (1), then there are cases in which an intrusion is *wanted* by the claimant, and yet there is a violation of the right to privacy. Consider Apology:

Apology

Person A has hurt the feelings of person C. Person A is truly regretful and wishes to give C a heartfelt apology. A is very nervous about giving the apology to C, and therefore, before giving the apology, A tells a close friend, B, how A wants to apologize to C. Unbeknownst to A, C eavesdrops on their conversation out of vengeance, in the hope of gaining knowledge of A’s personal information so she can tell others about it. C tells A that she has heard the apology, and A is truly relieved that she no longer has to deliver the apology face-to-face to C.

In this thought experiment, it seems that the intrusion is indeed wanted, since person A, had she been asked beforehand, would have wished that C would eavesdrop. But, C still clearly violates A’s (and possibly B’s) right to privacy. If this is correct, it demonstrates that it cannot be a necessary condition that the access is *unwanted* by the claimant, in order for a privacy violation to occur. For that reason, both control theorists and access theorists must accept that there are cases where the right to privacy is violated by an access that is, at least to some extent,

wanted by the claimant.¹⁹ In order to recognize this, the control theorist and the access theorists could simply exclude ‘unwanted’ from their respective accounts. The definitions would then read:

CA3: For any agent A to have her right to privacy violated, there is a necessary and sufficient condition that must be satisfied: Agent A has involuntarily lost Negative Control over the access to personal information P about agent A.

AA3: For any agent A to have her right to privacy violated there are two each necessary and jointly sufficient conditions that must be satisfied: (1) Agent A has involuntarily lost Negative Control over the access to personal information P about A, and (2) agent B actually accesses P.

Before turning to the test case, let us first discuss the issue of who, or what, can cause a loss of control of the kind that is relevant for a violation of the right to privacy to occur.

The Loss of Control

Kevin Macnish has recently argued that the descriptive AA is preferable to the descriptive CA. He writes:

I argue that the control account does not capture significant aspects of what is meant by privacy, demonstrating that privacy and control can come apart. Hence control is neither necessary nor sufficient for privacy. By contrast, privacy and access do not come apart. As such, I hold that the access account is preferable to the control account. (Macnish 2018, p. 1)

However, as we saw in Macnish’s diary case in ‘[The Control Account of the Right to Privacy](#)’ section, he also talks of *violations*. He claims that there is no violation in the diary example, even though there is a loss of control. Thus, he seems to prefer the normative AA over the normative CA.²⁰

The diary case demonstrates that control cannot be a necessary and sufficient condition for a privacy violation, since control is lost in that example while no violation has occurred. We agree that no violation occurs in the diary case. However, we will argue that this is not due to a lack of access. Rather, it is due to the fact that the loss of control is the claimant’s own fault, since he forgot the diary on the table. To see this more clearly, consider another example: you are walking outside in a storm with your diary in your bag. Unfortunately, you forgot to zip the bag completely, so

¹⁹ Thanks to Beate Rössler for pointing out the following to us: what is wanted by A in Apology is not the intrusion itself, but to give C the apology. But then let us change the example so that A wants C to intrude, because then A would feel that they were even, and that A no longer had to feel bad about what she did to C. Or, change it so that A has voyeuristic tendencies and likes to be watched or listened to by others. In these cases, A’s right to privacy would be violated (a right is not automatically waived just because the claimant likes that others occasionally violates the right), and yet the intrusion would be wanted.

²⁰ If Macnish did not intend this to be a discussion of privacy *rights*, he should have made that more explicit, and probably abstained from using the word ‘violation’.

the wind blows your diary out of the bag. It lands on the sidewalk with the pages facing up. Another pedestrian is kind enough to pick it up for you, but as he does so, he cannot avoid reading some of the content. In this case, there is clearly no violation, even though someone gets access to information in the diary, while there is a loss of control. This shows that the lack of access itself does not explain the lack of violation in Macnish's diary example. What explains the lack of violation, is the fact that the loss of control is not due to the action(s) of another agent, of which that agent is responsible.²¹

Adam Moore has an example which can be used to demonstrate that the loss of control must be due to the action(s) of another agent in order for a violation of the right to privacy to occur.²² Moore's example is this:

The Accidentally Amplified Quiet Fight: A married couple, X and Y, are having another quiet fight behind closed doors. But this time an unanticipated gust of wind sweeps through the house, knocking down the front door, carrying and amplifying the couple's voices so that Stuart, who is washing his car in his driveway across the street, hears at least some of what X and Y have been saying.

In the accidentally amplified quiet fight case the right to privacy is not waived and it also appears not to be violated. (Moore 2003, p. 423)

Although X and Y have lost control over the access to the information, and the information has indeed been accessed, no violation of the right to privacy has occurred, according to Moore. The loss of control, and Stuart's access to the information, is merely due to an accident, and for that reason, no violation has occurred. And, since no violation would have occurred if X and Y had given Stuart access voluntarily, it seems that the access must be due to the action(s) of another agent in order for a violation to occur. This is of interest for at least two reasons: (1) given how much work the diary example does for Macnish, it is problematic for him if it turns out that it is not the absence of access that explains the absence of a violation, and (2) it suggests a new adjustment of both definitions. The adjustment consists in adding that the loss of control must be due to the action(s) of another agent, of which that agent is responsible. The definitions then read:

CA4: For any agent A to have her right to privacy violated, there is a necessary and sufficient condition that must be satisfied: Agent A has involuntarily lost Negative Control over the access to personal information P about agent A, *due to action(s) of agent B, of which B is responsible.*

AA4: For any agent A to have her right to privacy violated there are two each necessary and jointly sufficient conditions that must be satisfied: (1) Agent A has involuntarily lost Negative Control over the access to personal information

²¹ This does not mean that no violations will occur downstream. For example, publicizing the forgotten diary on the Internet would still constitute a violation. See Moore (2018) for a discussion on issues of forfeiting and waiving rights.

²² Moore gets this example from Rickless (2007).

P about A, *due to the action(s) of agent B, of which B is responsible*, and (2) agent B (or someone else) actually accesses P.

A Test Case

Let us now consider a test case to see which of the improved accounts best explains the violations that occurs in this case. Call the test case Wiretapping:

Wiretapping

Smith and Jones are neighbors. Unbeknownst to Jones, Smith wiretaps Jones's telephone, using a fancy device which allows Smith to listen in on Jones's conversations without violating Jones's property rights. As it happens, Jones is on vacation for several months, and therefore does not use the telephone in that time period.

Our intuition is that Smith clearly violates Jones's right to privacy in Wiretapping. But which account best explains this violation? Let us first consider the improved version of the CA. According to CA4, it is a necessary and sufficient condition that Smith has lost negative control over the access to information, and that this loss of control was due to the action(s) of another agent, of which that agent is responsible. This seems satisfied in Wiretapping. Jones has lost negative control over the access, since Smith can now listen to Jones's telephone conversations. And, this loss of control was due to action(s) of Jones, for which Jones was responsible, since he was the one who chose to wiretap Smith's phone.

What about the AA? According to AA4, it is a necessary condition that Smith *actually accesses* Jones's information. But in Wiretapping, it seems that Smith does *not* access information about Jones, since Jones does not use the telephone. It could be argued that Smith does in fact access some information about Jones, namely the information that Jones did not use the particular telephone in that particular period. We grant that Smith has access to this information. But we find it hard to see that the access to *that* information alone is what drives the strong intuition that Jones's right to privacy is violated by Smith. Even if the wiretap had randomly malfunctioned unbeknownst to Smith, so Smith did not get access to the information that Jones did not use the telephone, Smith would clearly still have violated Jones's right to privacy. This counts against the AA, since it is too narrow to account for the violation in Wiretapping.

Wiretapping shows that, pace the access theorists' arguments, access is not a necessary condition for a violation of the right to privacy. Moreover, since there would be no violation if Jones had voluntarily given Smith access, it cannot be a sufficient condition either. This is a genuine problem for the access theorists, and a problem that we do not see how they can escape by simply adjusting the definition of the AA.

The access theorist might object that we are stacking the deck of cards in favor of the CA. After all, since there is no actual access in Wiretapping, it is not surprising that the AA cannot account for the violation. Our response to this objection is that none of the examples or thought experiments provided by the access theorists in the literature so far seem to count decisively in favor of the AA, like Wiretapping

counts decisively in favor of the CA. We cannot think of an example, which stacks the deck of cards in favor of the AA, so we invite the access theorists to provide such an example. A possible candidate for such an example is Judith Jarvis Thomson's seminal X-ray case:

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. (Thomson 1975, p. 304)

Thomson points out that your right to privacy has not been violated just because you no longer have control over whether your neighbor looks through your wall or not. It would only be violated, when the neighbor *actually*²³ trains the X-ray device on the wall (Thomson 1975, p. 305). Access theorists often turn to the X-ray case in order to show why the AA is preferable to the CA. We will show that the improved version of the CA can easily handle the X-ray case.

Let us first compare Wiretapping to the X-ray case. We agree that there is no violation in the X-ray case, unless the neighbor actually trains the X-ray on the wall. It might seem, *prima facie*, that on the CA4, there is a violation in the X-ray case, since control is lost due to the neighbor's actions (the invention of the X-ray device). But recall that the relevant form of control on the CA4 is Negative Control. In order for Negative Control to be lost, someone must *attempt* to get access, and in Thomson's case, the neighbor does *not* attempt to get access. To see clearly how this is an effective rejoinder to Thomson, let us return to the distinction between Negative Control and Republican Control which we introduced in an earlier section.

Republican Control is lost simply by virtue of someone else having the ability to access your information. They do not need to use this ability.²⁴ In Thomson's case, Republican Control is lost when the neighbor invents the X-ray device, but Negative Control is not lost. In Wiretapping, on the other hand, someone tries to get access, so Negative Control is lost. Thus, Thomson's attempt to make a *reductio* on the CA does not cut any ice against CA4.²⁵ Note also that Macnish's diary example does cut any ice against the CA4 either, since the loss of control in this example is also a loss of Republican Control, not a loss of Negative Control.

It seems that when we compare the improved versions of the two accounts, we have at least a *pro tanto* reason to prefer the CA over the AA. Only the CA can explain the violation in Wiretapping. This does not mean, however, that the CA is

²³ Note that this counts in favor of our earlier point that the access must be *actual* access, not only the *ability* to access.

²⁴ In 'The Access Account of the Right to Privacy' section, we argued that if the access in the AA is the *ability* to access, it would collapse into a *type* of CA. The type of CA it would collapse into is a republican CA.

²⁵ The distinction between Negative Control and Republican Control saves the control theorists from several objections in which the access theorists seem to think that a loss of Republican Control must be a violation on the CA. This shows the importance of specifying that the CA should only be concerned with losses of Negative Control.

preferable to the AA, all things considered. It might be that there are other problems with either of these accounts, which need to be accounted for, and that doing so reveals that in fact the AA comes out on top.

Concluding Remarks

In this paper, we have offered several ways in which both the control account and the access account of the right to privacy can be improved. We then tested the improved versions of the accounts to see which of them best explains the violation in Wiretapping. It turned out that the CA could explain the violation, while the AA could not. This gives us a pro tanto reason to favor the CA over the AA.

In the introduction, we claimed, following Kevin Macnish, that the discussion about which account of the right to privacy is the correct one is of tremendous importance for our normative evaluations of state surveillance. For example, when discussing the potential wrongdoing associated with the NSA's collection of data about people, and the Edward Snowden's subsequent whistleblowing, a lot hangs on whether the CA or the AA is correct. Macnish argued that if the CA is correct, then the NSA is violating citizens' right to privacy, but if the AA is correct, there is no such violation. This remains true with the adjustments we have suggested for the two definitions. On the CA4, the NSA's violation consists in a loss of negative control, by undermining people's ability to prevent the NSA (and others) from getting actual access to the information. When the information is stored in the NSA's database, the NSA has definitively undermined people's ability to control the access to the information, even if no employee of the NSA (or others) ever looks at the information. On the AA4, no violation occurs until an employee actually looks at the information.

What we have argued in this paper does not resolve the dispute between the control theorists and the access theorists decisively. But, if we are correct, then there is a pro tanto reason for saying that many instances of surveillance do in fact constitute violations of the right to privacy, even when the information in question is not actually accessed. As with any pro tanto reason, this one may be overruled by other reasons.

References

- Allen, Anita. 1988. *Uneasy Access: Privacy for Women in a Free Society*. Lanham, MD: Rowman & Littlefield Publishers.
- Allen, Anita. 1999. Coercing Privacy. *William and Mary Law Review* 40(3): 723–757.
- Allen, Anita. 2003. *Why Privacy isn't Everything: Feminist Reflections of Personal Accountability*. Lanham, MD: Rowman & Littlefield Publishers.
- Altman, Irwin. 1976. Privacy: A Conceptual Analysis. *Environment and Behavior* 8(1): 141.
- Berlin, Isaiah. 1969. *Two Concepts of Liberty*. Oxford: Clarendon Press.
- Bezanson, Randall P. 1992. The Right to Privacy Revisited: Privacy, News, and Social Change, 1890–1990. *Northwestern University* 80(5): 1133–1175.
- Bok, Sissela. 1989. *Secrets: On the Ethics of Concealment and Revelation*. New York: Vintage Books.

- Farber, Daniel A. 1993. Book Review: Privacy, Intimacy, and Isolation by Julie C. Inness. *Constitutional Commentary* 10: 510–519.
- Feinberg, Joel. 1985. *Offence to Others*. Oxford: Oxford University Press.
- Fried, Charles. 1968. Privacy. *Yale Law Journal* 77(3): 475–493.
- Gavison, Ruth. 1980. Privacy and the Limits of Law. *Yale Law Journal Article* 89(3): 421–471.
- Goldberg, Ian, Austin Hill, and Adam Shostack. 2001. Trust, Ethics and Privacy. *Boston University Law Review* 81(2): 407–422.
- Gross, Hyman. 1971. Privacy and Autonomy. In *Nomos XIII: Privacy*, pp. 169–181.
- Hoye, Matthew, and Jeffrey Monaghan. 2015. Surveillance, Freedom and the Republic. *European Journal of Political Theory* 17(3): 343–363.
- Inness, Julie C. 1992. *Privacy, Intimacy, and Isolation*. Oxford: Oxford University Press.
- Macnish, Kevin. 2018. Government Surveillance and Why Defining Privacy Matters in a Post-Snowden World. *Journal of Applied Philosophy* 35(2): 417–432.
- Margulis, Stephen T. 1977. Conceptions of Privacy: Current Status and Next Steps. *Journal of Social Issues* 33(3): 5–21.
- Miller, Arthur R. 1971. *The Assault on Privacy: Computers, Data Banks, and Dossiers*. Ann Arbor: University of Michigan Press.
- Moore, Adam D. 2003. Privacy: Its Meaning and Value. *American Philosophical Quarterly* 40(3): 215–227.
- Moore, Adam. 2008. Defining Privacy. *Journal of Social Philosophy* 39(3): 411–428. <https://doi.org/10.1111/j.1467-9833.2008.00433.x>.
- Moore, Adam D. 2010. *Privacy Rights: Moral and Legal Foundations*. University Park: Pennsylvania State University Press.
- Moore, Adam D. 2018. Privacy, Interests, and Inalienable Rights. *Moral Philosophy and Politics* 5(2): 327–355.
- Newell, Bryce Clayton. 2018. Privacy as Antipower. In Pursuit of Non-Domination (Foreword). *European Data Protection Law Review* 4(1): 12–16.
- Parent, W. A. 1983. Privacy, Morality, and the Law. *Philosophy and Public Affairs* 12: 269–288.
- Parker, Richard. 1974. A Definition of Privacy. *Rutgers Law Review* 27: 275–296.
- Pettit, Philip. 1999. *Republicanism: A Theory of Freedom and Government*. Oxford: Oxford University Press.
- Reiman, Jeffrey H. 1995. Driving to the Panopticon: A Philosophical Exploration of the Risks to Privacy Posed by the Highway Technology of the Future. In *Privacy*, ed. Eric Barendt, 159–176. Dartmouth: Ashgate.
- Rickless, C. Samuel. 2007. The Right to Privacy Unveiled. *San Diego Law Review* 44(4): 809–846.
- Roberts, A. 2014. A Republican Account of the Value of Privacy. *European Journal of Political Theory* 14(3): 320–344.
- Rössler, Beate. 2005. *The Value of Privacy*. Cambridge: Polity Press.
- Rubel, Alan. 2011. The Particularized Judgment Account of Privacy. *Res Publica* 17(3): 275–290.
- Ryan, M., and M. R. Calo. 2010. The Boundaries of Privacy Harm. *Indiana Law Journal* 86: 1131.
- Scanlon, Thomas. 1975. Thomson on Privacy. *Philosophy and Public Affairs* 4(4): 315–322.
- Schoeman, Ferdinand. 1984. *Philosophical Dimensions of Privacy*. Cambridge: Cambridge University Press.
- Thomson, Judith Jarvis. 1975. The Right to Privacy. *Philosophy & Public Affairs* 4(4): 295–314.
- van den Haag, Ernst. 1971. On Privacy. In *Privacy: Nomos XIII*, pp. 149–168.
- van der Sloot, B. 2018. A New Approach to the Right to Privacy or How the European Court of Human Rights Embraced the Non-domination Principle. *Computer Law and Security Review* 34(3): 539–549.
- Warren, Samuel D., and Louis D. Brandeis. 1890. Right to Privacy. *Harvard Law Review* 4(5): 193–220.
- Westin, Alan F. 1970. *Privacy and Freedom*. London: Bodley Head.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Privacy Rights, and Why Negative Control is *not* a Dead End: A Reply to Munch and Lundgren¹

Abstract

Lauritz Munch and Björn Lundgren have recently replied to a paper published by us in this journal. In our original paper, we defended a novel version of the so-called ‘control theory’ of the moral right to privacy. We argued that control theorists should define ‘control’ as what we coined ‘Negative Control’. Munch and Lundgren have recently provided a range of interesting and challenging objections to our view. Independently of each other, they give almost identical counterexamples to our definition of Negative Control. In this comment, we show that while the counterexamples are genuine counterexamples, they do not force us to abandon the idea of Negative Control. Furthermore, we reply to two additional objections raised by Lundgren. One of these replies involves giving a new account of what the relation is between the concept of privacy, and the right to privacy.

Introduction

In this journal, we have recently defended a novel version of the so-called ‘control theory’ of the moral right to privacy (Mainz & Uhrenfeldt 2020). Lauritz Munch and Björn Lundgren have independently of each other replied to our paper with a range of interesting and challenging objections (Munch 2021; Lundgren 2021). In this comment, we reply to the most important ones.

In our original paper, we tried to show why there is at least a *pro tanto* reason to favor the control theory over the rival ‘access theory’. Classic versions of the control theory of the moral right to privacy hold, roughly, that an agent A’s right to privacy is violated if, and only if, A does not have the relevant type of control over the access to A’s personal information. The version of the rival access theory that we discussed in our original paper adds a necessary condition for A’s right to privacy to be violated; that agent B actually accesses A’s personal information.

One of the crucial features of our version of the control theory is that it specifies how the control theorist should define the term ‘control’. We argued that the control theorist should define control as what we coined ‘Negative Control’. Based on the three well-known types of freedom - negative freedom, positive freedom, and republican freedom - we formulated three corresponding types of control. To wit, we contrasted Negative Control with Positive Control and Republican Control, respectively.² By defining control as Negative Control, we argued, the control theorist can avoid all the classic objections to the control theory. The reason for this is that all of the classic objections to the control theory assume a definition of control that is either Positive Control or Republican Control. On our account, agent A’s right to privacy is violated, if, and only if, A involuntarily loses Negative

¹ We thank Lauritz Munch, Frej Klem Thomsen, Jens Damgaard Thaysen, and Jørn Sønderholm, for useful comments on an earlier version of this paper.

² We defined Positive Control like this: Agent A enjoys Positive Control over the access to relevant information P, if, and only if, A tries (or could try) to give agent B actual access to P, and succeeds. And, we defined Republican Control like this: Agent A enjoys Republican Control if, and only if, agent B does not have the ability to get access to relevant information P about A. (Mainz & Uhrenfeldt 2020, 7).

Control due to the actions of agent B, for which B is responsible (Mainz & Uhrenfeldt 2020, 12). We defined Negative Control as follows:

Negative Control: Agent A enjoys Negative Control over access to relevant information P, if, and only if, A is capable of preventing agent B, who attempts to access, from accessing P. (Mainz & Uhrenfeldt 2020, 7)³

Munch and Lundgren provide almost identical counterexamples to the definition of Negative Control. In light of these counterexamples, we suggest how the definition could be altered. The alteration involves incorporating some of the components of a more recent version of Negative Control that Mainz has recently put forward in this journal (Mainz, *forthcoming*).⁴

In the next section, we discuss how the definition of Negative Control can be altered in order to accommodate the counterexamples provided by Munch and Lundgren, respectively. In the final section, we reply to two additional objections raised by Lundgren.

Two Counterexamples to Our Account

Let us begin with two almost identical counterexamples offered by Munch and Lundgren, respectively. These counterexamples purport to show that a loss of Negative Control is not a necessary condition for a violation of the right to privacy. To illustrate this point, Munch provides a hypothetical that is a modified version of one of our hypotheticals. Our original hypothetical is the following:

Wiretapping Smith and Jones are neighbors. Unbeknownst to Jones, Smith wiretaps Jones's telephone, using a fancy device which allows Smith to listen in on Jones's conversations without violating Jones's property rights. As it happens, Jones is on vacation for several months, and therefore does not use the telephone in that time period (Mainz & Uhrenfeldt 2020, 13).

Wiretapping purports to show that the rival access theory cannot explain the intuition that Smith violates Jones' right to privacy, because Smith does not actually access Jones' personal information.⁵ The control theory, on the other hand, can easily explain this intuition, if control is defined as Negative

³ The paper is only available online and lacks pagination, so the page numbers refer to the pages in the online version, starting from 1.

⁴ (Mainz, *forthcoming*) was accepted for publication before Munch and Lundgren's replies were published. The version of Negative Control put forward in that paper was therefore not supposed to handle the objections from Munch and Lundgren.

⁵ As Munch says in footnote 8 in his reply, we might interpret Wiretapping such that Smith actually accesses at least some information about Jones - for instance the fact that Jones is not using the phone. But as we say in the paper, the verdict would be the same even if the wiretap randomly malfunctions so that Jones does not even get access to the information that Smith is not using the phone (Mainz & Uhrenfeldt 2020, 13).

Control: Smith attempts to access Jones' personal information, but Jones is not capable of preventing Smith from accessing. Now, Munch provides an altered version of Wiretapping, which he calls

Wiretapping #2. Smith and Jones are neighbors. Smith wiretaps Jones's telephone, using a fancy device which allows Smith to listen in on Jones's conversations without violating Jones's property rights. Unbeknownst to Smith, Jones has an even fancier device enabling him to both monitor the extent to which he is being subjected to wiretapping and shut down the tapping at the mere push of a button. Jones does *not*, however, deploy his device to prevent Smith's plan (Munch 2021, 5).

Wiretapping #2 is a counterexample to our definition of Negative Control, because it demonstrates that a violation of the right to privacy can occur, even when no one loses Negative Control. Jones is in fact capable of preventing Smith, who attempts to access, from accessing. He just decides not to make use of this capability. Even so, Smith violates Jones' right to privacy. Thus, a loss of Negative Control is not a necessary condition for a violation of the right to privacy.

Lundgren provides a counterexample that is almost identical to Munch's Wiretapping #2. He writes:

“Imagine a case in which Smith is prevented from accessing Jones's phone not because of a malfunctioning device, but because Jones has a machine to prevent wiretapping. In this case, Jones retains negative control of his private information. However, we may still want to claim—as in Mainz and Uhrenfeldt—that Smith has violated Jones's right to privacy.” (Lundgren 2021).

We grant that Smith violates Jones' right to privacy in the two hypotheticals, and we acknowledge that they are clear and cleverly constructed counterexamples.⁶ We do not, however, think that this leads us to the conclusion that the idea of Negative Control is a ‘dead end’, as the title of Munch's reply suggests. Rather, we think that the definition of Negative Control can be altered to handle the counterexamples, without abandoning the underlying idea that control should be interpreted as something akin to the idea of negative freedom.

It lies beyond the scope of this reply paper to provide a fully developed alternative to our original definition of Negative Control. However, the version of Negative Control put forward in (Mainz, *forthcoming*) contains elements that can work as a useful starting point. For present purposes, let us call this version

⁶ We do think, however, that the counterexamples are underspecified in an important sense. They say nothing about why Jones might choose not to push the button. Suppose that Jones decides not to push the button, because doing so would be extremely costly for him, or because he simply “freezes” in the situation. Now compare a situation in which Jones chooses not to push the button because he would actually like Smith to listen in on his conversations. We think that any plausible theory of rights should be able to say that there is a rights-violation in both cases. But it seems that the wrongness involved in the two cases are not identical. The wrongness that occurs in the former case seems much worse than the one that occurs in the latter. Nevertheless, we grant that a rights-violation occurs in both cases. As Munch points out, denying this would be akin to denying that an assaulter violates the rights of the assaultee, even if the assaultee is capable of fending off the assaulter (Munch 2021, 6).

Negative Control #2: An individual A has Negative Control over relevant information f with respect to B, if, and only if,

i) B does not attempt to access f (or attempts to give others access), or ii) B does attempt to access f (or attempts to give others access), but fails due to A's intentional actions directed at preventing B from accessing f , or, due to random circumstances, or, due to the incompetence of B,

and,

iii) A does not voluntarily let B access f .⁷ (Mainz, *forthcoming*).

Let us briefly clarify what Negative Control #2 holds. Negative Control #2 implies that A has privacy *iff* either i) or ii) is satisfied, while iii) is also satisfied. Correspondingly, A does not have privacy if neither i) nor ii) are satisfied, if iii) is not satisfied, or if neither i), ii), or iii) are satisfied.

Importantly, Mainz did not *defend* this definition in (Mainz, *forthcoming*). He used it merely to show what the Negative Control account might look like if it was used to define the *concept* of privacy. Negative Control #2 was thus not intended as a definition of the type of control that is at stake in the right to privacy. Nevertheless, parts of it can be used to avoid the counterexamples from Munch and Lundgren. Here is how.

Let us first consider condition ii). Thanks to condition ii), the definition avoids the two counterexamples, because Smith does not fail his attempt to access Jones' personal information - let alone fail because of any of the reasons described in ii). Because Smith does not fail his attempt, Jones does not have Negative Control, and thus Smith violates Jones' right to privacy. Admittedly, it seems *prima facie* strange to say that Jones does not have control even though he decides not to deploy the device. We contend, however, that this is only superficially problematic. By analogy, consider how we normally think about property rights. If we have a property right in a painting, then we have - *inter alia* - a control right over the painting. Nevertheless, this control right is plausibly violated when we decide not to fend off a burglar who is trying to steal the painting. This is so even if we are perfectly capable of fending off the burglar.⁸

⁷ Note that the information is called f in this definition, while it was called P in the original definition from (Mainz & Uhrenfeldt 2020). Note also that while the original definition concerns *agents* in general, this definition is concerned with *individuals*.

⁸ One difference to note between Wiretapping #2 and Lundgren's counterexample is that the latter does not explicitly state whether Jones deploys the device, while the former says explicitly that Jones does *not* deploy the device. But given that Lundgren stipulates that Smith is *prevented* from accessing Jones' phone, it seems that Jones deploys the device. However, regardless of whether Jones deploys the device or not, ii) can elegantly handle the example. Jones either deploys the device or he does not. If he *does*, then Smith does not violate Jones' right to privacy because Smith's attempt to access fails because of Jones' intentional actions directed at preventing Smith from accessing. In that case, Jones still enjoys Negative Control, and Smith does plausibly not violate Jones' right to privacy. If Jones does *not* deploy the device, then Smith does not fail his attempt to access Jones' personal information - let alone fail because of any of the reasons described in ii). Because Smith does not fail his attempt, Jones loses Negative Control, and thus Smith violates Jones' right to privacy.

Had Jones deployed his fancy device and jammed Smith's wiretap, then Smith's attempt to access would have failed because of Jones' intentional actions directed at preventing Smith from accessing. In that case, Jones would still enjoy Negative Control, and Smith would plausibly not have violated Jones' right to privacy.

What about condition i)? Condition i) does not help us avoid any of the two counterexamples. It does, however, provide a reply to another of Munch's objections. As Munch notes in footnote 3 in his reply, our original definition of Negative Control implies that A *only* has Negative Control in the moment where someone actually attempts to access the information. Strangely, A does not have Negative Control when *no one* attempts to access. Condition i) lets us escape this admittedly strange implication of our original definition. The reason is that i) explicitly states that no one attempts to access. So, given that we have a disjunction consisting of i) and ii), it is sufficient for having Negative Control that no one attempts to access one's personal information.

Now, what about condition iii)? We suggest that this condition should be dropped. The reason is that condition iii) is implausible when we are concerned with the *right* to privacy, because including iii) implies that Smith violates Jones' right to privacy, if Jones voluntarily tells Smith a personal secret about himself. This would be a very unfortunate result, so we must drop condition iii).⁹

Negative Control #2 constitutes a promising starting point for developing a plausible definition of what kind of control is at stake in the control theory of the moral right to privacy. Moreover, it straightforwardly avoids Munch and Lundgren's counterexamples without abandoning the core idea of Negative Control.

Two Further Objections from Lundgren

Let us now move on to two additional objections raised by Lundgren. The first objection is that we do not recognize that privacy is the object of the right to privacy. By defining the right to privacy in terms of control, Lundgren says, one must also define the concept of privacy in terms of control (Lundgren 2021, 3). Lundgren has recently defended this view of the relation between the right to privacy and the concept of privacy thoroughly in (Lundgren 2020). This contribution to the literature is very welcome, and it opens up the underdeveloped discussion of what the relation is between the right to privacy and the concept of privacy. Lundgren claims that because we define the *right* to privacy in terms of control, we must subscribe to a control-based definition of the *concept* of privacy. This is a problem for us, Lundgren says, because we ignore the counterexamples to theories that define the concept of privacy in terms of control.

We do not think that Lundgren gets the relation between the right to privacy and the concept of privacy completely right, and - consequently - we think that his objection to us misfires. We grant Lundgren's point that privacy is the object of the right to privacy. The right to privacy is a right to be in a condition of *privacy*. However, it is a *non sequitur* to say that by endorsing a control-based theory of the *right* to privacy, we are therefore necessarily committed to endorsing a control-based theory of the *concept* of privacy. To see why this is a *non sequitur*, consider the difference between claiming that

⁹ For discussion of similar cases, see the Too Much Info cases in (Mainz & Uhrenfeldt 2020).

“the right to privacy is a right to be in a condition of privacy[*defined in terms of control*]”,

and claiming that

“the right[*defined in terms of control*] to privacy is a right to be in a condition of privacy”.

In the first claim, the control-part attaches to the *concept* of privacy, while in the second claim, it attaches to the *right* to privacy. We agree with Lundgren that the concept of privacy should *not* be defined in terms of control. In fact, Mainz has recently defended the view that privacy should be defined in terms of access (Mainz, *forthcoming*). Lundgren does not seem to recognize – neither in (Lundgren 2020), nor in his reply to us – the difference between the two claims above. The difference between the two claims allows us to reject Lundgren’s view, because it allows us to reject the view that if the right to privacy should be defined in terms of control, then so should the concept of privacy. Simply put, we can consistently hold that we have control rights (whatever that means exactly) over the access to our personal information (whatever that means exactly).

An analogy to property rights may be helpful here: It is one thing to be in a condition of possessing a car, and another thing to have a property right – which conventionally includes a *control* right – over the car. A car thief is in a condition of possessing the car, but he does not have a control right over the car. And, the owner of the stolen car has a control right over the car, but he is not in a condition of possessing the car (because the car thief is). We can consistently endorse the view that property rights should be defined in terms of control (among other things), while also endorsing the view that the concept of possession should not. Still, having a property right in X is to have a right to possess X.

Something similar holds for the relation between the concept of privacy and the right to privacy. We can define the right to privacy in terms of control, without being forced to define the concept of privacy in terms of control. This is consistent with the view that the concept of privacy is the object of the right to privacy. It is not clear why the relation we have sketched out here does not constitute the ‘appropriate consistency’ between the definitions of the right to privacy and the concept of privacy that Lundgren is asking for (Lundgren 2021, 4).¹⁰

The second objection raised by Lundgren is that Wiretapping is a problematic test case for whether our version of the control theory is more plausible than the rival access theory. He thinks that Wiretapping does nothing to convince someone who does not already have control-based intuitions (Lundgren 2021, 4). Lundgren thinks that there are two viable options for an access theorist to reply to Wiretapping. The first one is simply to deny that one shares the intuition that Smith violates Jones’ right to privacy. The second one is to reformulate the access theory, such that it can accommodate the intuition that Smith violates Jones’ right to privacy.

Regarding the first option, we contend that very few people would be willing to bite the bullet and say that Smith does not violate Jones’ right to privacy in Wiretapping. The reason why we used this

¹⁰ Keep in mind that we are *not* arguing that there is no relation between the concept of privacy, and the right to privacy. A loss of the former is indeed a necessary condition for a violation of the latter. All we are saying is that accepting the view that the right to privacy is a control right does not force us to accept the view that the concept of privacy should be defined in terms of control.

exact case is that it seems to be a paradigmatic example of a privacy violation. The methodological motivation for choosing the example is that if we are trying to reach a reflective equilibrium between the considered judgment about Wiretapping, and a general theory of the right to privacy, then we think that the general theory of the right to privacy has to give, until it is consistent with the considered judgment that Smith at least wrongs Jones in Wiretapping. In other words, if the pre-theoretical intuition is sufficiently strong in this case, then the intuition should guide us in our theory construction. This is what we mean when we say that Wiretapping is a test case.¹¹

Regarding the second option, Lundgren claims that the access theorist can “[...] easily agree with the intuition that Smith has violated Jones’ privacy, but deny that the limited access conception cannot explain this.” (Lundgren 2021, 4). The access theorist can agree with the intuition, if she drops the view that actual access to private information is a necessary condition for a violation of the right to privacy:

“[...] the right to privacy protects against substantial risks of access, not merely actual access. That is, while actual access to someone’s private information might be a necessary criterion for when someone’s privacy is diminished, it is not clear that we should hold that actual access is a necessary criterion for when the right to privacy is violated.” (Lundgren 2021, 4).

This is a view that Lundgren defends in (Lundgren, *forthcoming*). Let us call it the ‘substantial risk view’. Notice that Lundgren decides to appeal to a view that has - as Lundgren admits - never been defended in print, in order to accommodate Wiretapping.¹² Given this, it does not appear to be an ‘easy’ concession on behalf of access theorists. If this is the best option - and Lundgren seems to think that it is - then it is worth highlighting that the solution has taken him quite far away from the original access theories which crucially hold that actual access is a necessary condition for a violation of the right to privacy. So, even if Lundgren’s new version of the access theory turns out to be correct, then our Wiretapping case still has bite against the original access theories.¹³

However, we believe that there are at least two reasons why the substantial risk view that Lundgren appeals to is bound to fail. The first reason is the following: Recall Lundgren’s counterexample regarding a device that can block wiretapping. Now suppose that there is a substantial risk that the device will malfunction, but luckily for Jones, the device works. Jones deploys the device before Smith gets a chance to listen in on Jones’ conversations. On Lundgren’s substantial risk view, Smith violates Jones’ right to privacy. On our view, Smith does *not* violate Jones’ right to privacy, because Smith fails due to Jones’ intentional actions directed at preventing Smith from accessing (Jones has Negative Control). To see why our verdict of this case is more plausible than Lundgren’s, consider a brief example: You see that your neighbor is about to peep into your bedroom through the window. Before he gets a chance to look, you close your curtains. The curtains are old, so there is a substantial risk that they will fall down when you close them. Luckily, they do not fall down, and you

¹¹ And, as we say, if the access theorist feels that we are stacking the deck of cards in favor of the control theory here, then we invite the access theorist to provide an example that stacks the deck of cards in favor of the access theory (Mainz & Uhrenfeldt 2020, 13). The counterexamples from Munch and Lundgren may in fact be just such an example.

¹² It is not completely true that no one has defended this view in print before. Munch has defended something very similar in (Munch 2020).

¹³ See e.g. (Thomson 1975). Thomson is probably the most prominent access theorist, and she explicitly defends the view that actual access is a necessary condition for a violation of the right to privacy (Thomson 1975, 304).

successfully block your neighbor's attempt to peep into your bedroom. In this case, it seems more intuitive to say that there was a morally problematic attempt to violate your right to privacy, but that the attempt fails.¹⁴ This is what our view holds, while Lundgren's does not. Lundgren's view seems more akin to conceptually categorizing attempted murder as a successful murder.

The second reason takes the form of a *tu quoque*: If Lundgren opts for the substantial risk view, then he is faced with the same difficulty that he argues that our view is faced with - namely, that there is seemingly no 'appropriate consistency' between the *right* to privacy and the *concept* of privacy. On the substantial risk view, actual access is not a necessary condition for a violation of the right to privacy. And it is, for obvious reasons, not a sufficient condition either.¹⁵ If it is neither a necessary, nor a sufficient condition, then it is difficult to see how there is *any* relevant relation between the right to privacy and the concept of privacy. This is, as mentioned, the problem that Lundgren initially raised against *our* position.

Where does all of this leave our Negative Control account? When it comes to the right to privacy, Munch and Lundgren have convincingly shown that the original definition of Negative Control was flawed. In this comment, we have suggested that what we call Negative Control #2 can avoid their objections to our original definition. We think that Negative Control #2 constitutes a promising starting point for developing a more refined version of the Control Theory of the moral right to privacy. We leave it for another occasion to further develop such a theory.

References

- Lundgren, B. 2020. "A Dilemma for Privacy as Control." *Journal of Ethics*. 24: 165–175.
<https://doi.org/10.1007/s10892-019-09316-z>
- Lundgren, B. 2021. "Confusion and the Role of Intuitions in the Debate on the Conception of the Right to Privacy". *Res Publica*. <https://doi.org/10.1007/s11158-020-09495-9>
- Lundgren, B. "How we can make sense of control-based intuitions for limited access-conceptions of the right to privacy." *Journal of Ethics and Social Philosophy*. Forthcoming.
- Mainz, J. Uhrenfeldt, R. 2020. "Too Much Info: Data Surveillance and Reasons to Favor the Control Account of the Right to Privacy." *Res Publica*. <https://doi.org/10.1007/s11158-020-09473-1>.
- Mainz, J. "An Indirect Argument for the Access Theory of the Right to Privacy." *Res Publica*.

¹⁴ When we wrote the original paper, we were not fully convinced that actual access is a necessary condition for a violation of the right to privacy. For this reason, we could consistently hold at the time that Smith violates Jones' right to privacy in Wiretapping, even when Smith does not get access to any personal information about Jones. Since then, we have come around to the view that a loss of privacy *is* a necessary condition for a violation of the right to privacy, and therefore, we no longer believe that Smith violates Jones' *right* to privacy. Luckily, this does not force us to abandon the strong intuition that Smith is wronging Jones. We can simply say - as Lundgren suggests (Lundgren 2021, 5) - that Smith merely *attempts* but fails to violate Jones' right to privacy, and that this attempt is wrongful.

¹⁵ Actual access as a sufficient condition is a non-starter. This view implies that if you voluntarily give out personal information, then *your* right to privacy is violated.

Forthcoming.

Munch, L. 2020. “The Right to Privacy, Control Over Self- Presentation, and Subsequent Harm”.

Journal of Applied Philosophy. 37(1): 141-154.

Munch, L. 2021. “Why ‘Negative Control’ is a Dead End: A Reply to Mainz & Uhrenfeldt.” *Res*

Publica.

Thomson, J. 1975. “The Right to Privacy.” *Philosophy & Public Affairs*. 4(4): 295–314.



An Indirect Argument for the Access Theory of Privacy

Jakob Mainz¹

Accepted: 18 June 2021

© The Author(s), under exclusive licence to Springer Nature B.V. 2021

Abstract

In this paper, I offer an indirect argument for the Access Theory of privacy. First, I develop a new version of the rival Control Theory that is immune to all the classic objections against it. Second, I show that this new version of the Control Theory collapses into the Access Theory. I call the new version the ‘Negative Control Account’. Roughly speaking, the classic Control Theory holds that you have privacy if, and only if, you can control whether other people know personal information about you. Critics of the Control Theory often give counterexamples, where privacy is either not diminished even though the claimant has lost control, or where privacy is diminished even though the claimant is in control. I argue that none of these alleged counterexamples work against the Negative Control Account. However, this is not a victory for the control theorist, because the Negative Control Account collapses into the Access Theory. The paper thus adds to the recent trend in the literature of favoring the Access Theory over the Control Theory.

Keywords Privacy · Control theory · Access theory · Personal information

Introduction

A significant part of the privacy literature consists of discussions about how best to define privacy. Among the prominent theories of privacy are the ones which David Matheson has called the ‘Limited Access Theory’,¹ the Narrow Ignorance Theory,²

¹ According to which ‘An individual A has informational privacy relative to another individual B and to a personal fact f about A if and only if there are extraordinary limitations on B’s ability to know f ’ (Matheson 2007, p. 253). This theory has been defended by prominent theorists like Ruth Gavison (1980) and Anita Allen (1988).

² According to which ‘An individual A has informational privacy relative to another individual B and to a personal fact f about A if and only if (1) f is undocumented and (2) B does not know f ’ (Matheson 2007, p. 253).

✉ Jakob Mainz
jtm@learning.aau.dk

¹ Aalborg University, Aalborg, Denmark

and the Broad Ignorance Theory.³ I will not rehearse the arguments for and against these theories in this paper, but note that at this stage of the discussion, many theorists seem to agree that these theories are essentially flawed. Most contemporary theorists seem to subscribe to some version of the so-called Control Theory instead.⁴ Although variations of the Control Theory seem to be the most popular ones, a host of alleged counterexamples has been raised against it.⁵ So far, the control theorists have not provided satisfactory replies to these counterexamples. In this paper, I try to provide a unified and effective reply to all of the most worrying counterexamples on their behalf. However, as we shall see, this reply implies that the Control Theory collapses into the so-called Access Theory. This is very problematic for the control theorist, given that access theories are generally considered the main rivals to the Control Theory (Macnish 2018).

Although most contemporary theorists subscribe to the Control Theory,⁶ this paper is part of a recent trend in the literature that suggests that the Access Theory is in fact superior to the Control Theory (Lundgren 2020; Macnish 2018). Björn Lundgren has recently argued that—due to what he calls the Parent/Macnish-dilemma—the Control Theory should be rejected, and the Access Theory should be endorsed instead (Lundgren 2020). However, Leonhard Menges has recently shown convincingly that the Parent/Macnish-dilemma can be resolved. His solution consists in interpreting ‘control’ as what he calls ‘source control’ (Menges 2020). As I will argue at the end of this paper, although Menges’s source control account is convincing in many regards, it should be rejected after all. The main contribution of this paper is twofold: First, I offer an alternative version of the Control Theory that is immune to the classic objections. So far, this is the only version of the Control Theory in the literature that achieves this. Second, I show how this new theory collapses into the Access Theory.

The paper is structured as follows: In the section ‘The Control Theory’, I introduce a definition of the Control Theory. In the section ‘Two Strategies for Refuting the Control Theory’, I explain two types of argumentative strategies that critics of the Control Theory have followed in order to refute it. I call these argumentative strategies the A-strategy and the B-strategy, respectively. In the section ‘Three Types of Control’, I introduce a distinction between three types of control: Negative Control, Positive Control, and Republican Control. In the section ‘Averting

³ According to which ‘An individual A has informational privacy relative to another individual B and to a personal fact *f* about A if and only if B does not know *f*’ (Matheson 2007, p. 259). This theory has been put forward by Matheson himself. Several theorists have recently endorsed altered versions of this theory. See e.g. (Blaauw 2013), (Kappel 2013), and (Fallis 2013).

⁴ Variations of the Control Theory can be found in (Warren and Brandeis 1890), (Westin 1970), (Fried 1968), (Rachels 1975), (Moore 2003; Moore 2010), (Gross 1971), (Parker 1974), (Matthews 2008), (Roessler 2005) (Benzanson 1991), (Goldberg, Hill, and Shostack 2001), (Altman 1976), (Calo 2011), (Miller and Weckert 2000), (Inness 1992), (Birnhack 2019), (Falls-Corbitt and McLain 1992), (Frey 2000), and (Froomkin 2000). Some of these theorists call it the ‘control account’ instead.

⁵ I write ‘alleged’ counterexamples, because I do not—for reasons I will spell out in this paper—believe that they are genuine counterexamples to the Control Theory. Throughout the paper, when I write ‘counterexample’ I mean an alleged counterexample, unless specified otherwise.

⁶ See (Menges 2020) for a recent defense of a novel version of the Control Theory.

Counterexamples', I argue that if 'control' is interpreted as Negative Control, and not Positive Control or Republican Control, then the control theorist can effectively avert both the counterexamples that follow the A-strategy and those that follow the B-strategy. In the section 'How Negative Control Collapses the Control Theory into the Access Theory', I argue that if the control theorist interprets control as Negative Control, then the Control Theory collapses into the rival Access Theory. In the section 'The Source Control Objection', I present Menges's source control account that avoids collapsing into the Access Theory. I show that even if this is true, the source control account should be rejected for independent reasons. In the final section, I make a few concluding remarks.

The Control Theory

In this section, I will introduce a definition of the Control Theory. This theory comes in many different variations, but central to all of them is—loosely—the idea that having privacy is a matter of having control. For the purposes of this paper, nothing of importance hangs on how exactly the Control Theory is spelled out, but I will follow David Matheson's semi-formalized version:

The Control Theory (CT)

An individual A has informational privacy relative to another individual B and to a personal fact f about A if and only if A controls whether B knows f . (Matheson 2007, p. 252).

It is helpful to note a few things about this definition of the CT. (I) The CT is non-normative. In itself, it says nothing about whether privacy is valuable, whether privacy rights exist, or what it takes to violate privacy rights if they do exist. (II) According to the CT, B must know f about A in order for A's privacy to be diminished relative to B and relative to f . Recent critics have pointed out that weaker epistemic relations than knowledge are sufficient for privacy to be diminished, and that the stronger the epistemic relation is, the more privacy is diminished (Blaauw 2013; Kappel 2013; Fallis 2013). I find this critique compelling, but I will bracket it for now, since it is fairly easy to see how a weaker epistemic relation can be replaced with 'knows' in the definition without turning it into something that is not a *control* theory. (III) The CT states a necessary and sufficient condition for A having privacy, namely that A controls whether B knows f about A.

Two Strategies for Refuting the Control Theory

Given (III), at least two effective strategies are available for a critic of the CT. Since the definition of the CT states a necessary *and* sufficient condition for privacy, a critic of the CT can attack the necessity-part, or she can attack the sufficiency-part.

Many objections to the CT take the form of a *reductio ad absurdum*, where a counterexample (often in the form of a thought experiment) is offered to show that control is either not necessary or not sufficient for privacy. The two types of strategies against the CT can thus be described in the following manner:

The A-strategy Show that privacy is sometimes diminished, even if control is *not* diminished.

The B-strategy Show that privacy is sometimes intact, even if control *is* diminished.

Counterexamples that follow the A-strategy aim to show that the CT is too narrow. That is, they aim to show that control is not sufficient for privacy. And counterexamples that follow the B-strategy aim to show that the CT is too broad. That is, they aim to show that control is not necessary for privacy. The reason why I frame the discussion in terms of the A-strategy and the B-strategy, and not just in terms of narrowness and broadness, is that the CT or variations of it have also been accused of being too narrow and too broad for reasons that are not related to control. For instance, some think that the CT is too narrow because it only concerns ‘informational privacy’ (Solove 2002). Others think that the CT is too broad, because not all ‘personal facts’ are private (*ibid.*).

Both strategies can be found in the literature, but the B-strategy seems to be the most common one in the works of prominent privacy scholars. Critics who follow the A-strategy often make use of variations of so-called ‘voluntarily divulgence cases’ (Parent 1983; Menges 2020). Critics who follow the B-strategy often make use of so-called ‘threatened loss cases’ (Parent 1983; Menges 2020), but as we shall see they also make use of others types of cases. In the next section, I will introduce three types of control, which will become crucial in the subsequent discussion.

Three Types of Control

Let me introduce a distinction between three types of control. Call them Negative Control, Positive Control, and Republican Control, respectively. As we shall see, these three types of control are inspired by the distinction between three types of freedom in the political philosophy literature:

Negative Control An individual A has Negative Control over relevant information f with respect to B, if, and only if,

- (i) B does not attempt to access f (or attempts to give others access), or (ii) B does attempt to access f (or attempts to give others access), but fails due to A’s intentional actions directed at preventing B from accessing f , or, due to random circumstances, or, due to the incompetence of B, and,

(iii) A does not voluntarily let B access f .⁷

Positive Control An individual A has Positive Control over relevant information f with respect to B, if, and only if,

(iv) A wants to give B access to f , and A can act so that B gets access to f .

Republican Control An individual A has Republican Control over relevant information f with respect to B, if, and only if,

(v) B could not get access to f if B tried.⁸

It is helpful to note a few things about these definitions. First, Negative Control is defined in a way that implies that A's privacy can be diminished in two ways. The first way to diminish privacy occurs if neither (i) nor (ii) are satisfied. The second way to diminish privacy occurs if (iii) is not satisfied.⁹ An obvious example of the first way to diminish privacy involves a peeping Tom who gets access to information about what A does in her bedroom by peeping in between the curtains. An obvious example of the second way to diminish privacy involves an exhibitionist A who wants peeping Tom to access the information about what A does in the bedroom, and therefore opens the curtains and lets Tom watch.¹⁰

Second, Positive Control is defined in a way that implies that A has it if A is able to give others access to f , regardless of whether or not they want to have access or not. An example of this involves an exhibitionist who forces others to look at her while she performs sexual acts.

Third, contrary to Negative Control, Republican Control is defined in a way that implies that A has it whether or not someone else attempts to access f . If A does not have Republican Control, then it follows that someone is able to access f , and this

⁷ To see why it is important to include the part about random circumstance, consider the following example: B is about to access A's personal information. A is not capable of preventing B from accessing, but just before B accesses, B is struck by lightning and dies. If the part about random circumstances were not included, it would follow from the definition that B diminishes A's privacy, which seems odd (thanks to Leonhard Menges for suggesting this example to me). Similarly, it is important to include the part about incompetence. Suppose, for example, that B attempts to peep in between A's curtains by jumping up and down on the sidewalk. But A lives on the 5th floor, so even when B jumps as high as he can, there is no chance that he will succeed. If the part about incompetence was left out, B would diminish A's privacy. Again, that would be odd. Thanks to an anonymous reviewer for pointing this out to me, and thanks to Jens Damgaard Thaysen for suggesting the example.

⁸ These definitions are revised versions of the definitions first put forward in (Mainz and Uhrenfeldt 2020).

⁹ In order to fully flesh out what Negative Control consists of, it would be necessary to explain what exactly constitutes an 'attempt' to access f . One might think, for example, that the mere fact that I attempt to get access to my neighbor's health records by reading his mind, does not count as a genuine attempt. Likewise, it would be necessary to explain if A loses Negative Control if the individual who attempts to access f fails for reasons unrelated to A's intentional actions directed at preventing the individual from accessing f . I will leave these and related questions for another occasion. Note, however, that if these questions give rise to counterexamples to the definition of Negative Control, then it is a problem for the control theorist, not for the argument I make in this paper.

¹⁰ One might think that the exhibitionist is in fact *exercising* control in this example. This may be true, but note that this form of control is *Positive Control*.

someone can have this control without ever getting close to f in any way, and without ever attempting to access f . An example of this involves a peeping Tom who is able to look through A's curtains if he wants to.

The distinctions between the three types of control loosely resemble the well-known distinctions between different types of freedom in the political philosophy literature. The distinction between Negative Control and Positive Control loosely resembles Isaiah Berlin's distinction between negative freedom and positive freedom (Berlin 1969). The notion of Republican Control loosely resembles Philip Pettit's notion of republican freedom (Pettit 1999).¹¹ What I mean by resemblance here is that the respective types of control at play in my distinctions between Negative Control, Positive Control, and Republican Control respectively, are relatively akin to the type of control that the claimant has if she has negative freedom, positive freedom, and republican freedom, respectively. The three types of control do not map on to the corresponding types of freedom perfectly. However, this is not a big problem. What matters is that interpreting control as Negative Control averts all the classic counterexamples to the CT, because all these counterexamples turn on interpretations of control that are either Positive Control or Republican Control.

Negative freedom is the absence of interference from others (Berlin 1969, pp. 15–22). Positive freedom, on the other hand, is the ability to do certain things (ibid., pp. 22–25). If, for example, someone stops you from running to wherever you want, then you do not have full negative freedom. If you, on the other hand, are physically disabled and unable to run, then you do not have full positive freedom.

Pettit thought that the combination of negative freedom and positive freedom does not capture all aspects of freedom. One can, Pettit thought, be unfree in an important way, even though one has full negative freedom and full positive freedom. A slave that is owned by a benevolent slave owner might have both full negative freedom and full positive freedom, and yet it seems strange to say that the slave is really free, since the slave owner is able to interfere at any given time if he so chooses (Pettit 1999, pp. 32–35).

Here is how the definitions of Negative Control, Positive Control, and Republican Control are related to the three types of control that the claimant has if she has negative freedom, positive freedom, and republican freedom, respectively: Just like negative freedom, Negative Control has to do with the lack of interference from others. Part (i) of the definition of Negative Control reflects this aspect of negative freedom. However, Berlin did not only think of negative freedom as a condition of a lack of interference from others. He also wrote: 'The defense of [negative] freedom consists in the "negative" goal of warding off interference' (Berlin 1969, p. 20). Thus, Berlin's idea of negative freedom also has an 'active' component; it also has to do with being able to 'ward off' interference. Part (ii) of the definition of Negative Control reflects this active of negative freedom. If one is not able to 'ward off' someone

¹¹ Others have explored the relation between privacy and republicanism. See e.g. (Newell 2018), (Roberts 2014), (van der Sloot 2018), and (Hoye and Monaghan 2018). However, these theorists do not apply the notion of republican freedom to the notion of control as I do, but rather argue that privacy is important for retaining republican freedom.

else's attempt to access the personal information in question, then one does not have Negative Control over the access to said information.¹² If one puts up certain irreversible obstacles for oneself, then—at least on one interpretation of the notion—one's negative freedom is diminished. Part (iii) of the definition of Negative Control reflects this aspect of negative freedom.¹³

Just like positive freedom, Positive Control has to do with being able to do certain things. If one wants to give someone access to a piece of personal information, and one is able to give this someone access, then one has Positive Control. If one is not able give this someone access to the information, then one does not have Positive Control over the access to said information.¹⁴

Just like republican freedom, Republican Control has to do with others not being able to do certain things to you. If others are not able to get access to one's personal information, then one has Republican Control. If others are able to get access to one's personal information, then one does not have Republican Control over the access to said information.¹⁵

I have now defined and explained what I mean by Negative Control, Positive Control, and Republican Control, respectively. In the next section, I will show that if the control theorist makes clear that she interprets control as Negative Control, then the counterexamples to the CT that follow the A-strategy, and those that follow the B-strategy, are averted. Henceforth, I will call the resulting theory of interpreting control as Negative Control, the 'Negative Control Account'.

Averting Counterexamples

In this section, I will show how the Negative Control Account averts the counterexamples that follow the A-strategy, and the counterexamples that follow the B-strategy. I believe that this goes for *all* counterexamples that follow one of these strategies. If the control theorist makes clear that control should be interpreted as Negative Control, then none of the counterexamples cut any ice against the CT. I will give two examples of this in relation to the A-strategy, and six examples of this in relation to the B-strategy, but I believe that the point generalizes. I will present

¹² Note that in this case, both i) and ii) will not be satisfied, and therefore A's privacy is diminished.

¹³ An anonymous reviewer suggested to me that it seems strange that the definition of Negative Control implies that A has Negative Control in a situation where A does not voluntarily let someone else access *f*, and no one attempts to access. The definition is formulated like this because Negative Control is something that you have under normal circumstances, when no one even attempts to interfere. Plausibly to my mind, it would be even stranger to suggest that A only has Negative Control when someone actually attempts to access. Presumably, in order for A to lose control, she must have it in the first place.

¹⁴ See the 'Too Much Info' cases from (Mainz and Uhrenfeldt 2020) for examples of this.

¹⁵ An anonymous reviewer suggested to me that if A has Positive Control (is able to share information with others if she wants to), then A lacks Republican Control. I do not see how that follows. Suppose that A wants to share information *f* with B, and she goes on and does so. Does this mean that A did not have Republican Control? No. In order for A to lack Republican Control, it must be the case that B could just get access to *f* anyway, even if A had not herself decided to share *f* with B. But it is perfectly possible that B cannot access *f* if she wants to, even if A voluntarily shares *f* with B.

the counterexamples one at a time. The point I make after each counterexample is exactly the same: If control is interpreted as Negative Control, then the counterexample in question loses its bite. I illustrate this point by repeating the same texts after each counterexample, replacing only the name of the author of the counterexample, and the description of the case in question.

Let us begin with the counterexamples that follow the A-strategy.¹⁶ The first counterexample that follows the A-strategy comes from William Parent:

All of these definitions [the control definitions of privacy] should be jettisoned. To see why, consider the example of a person who voluntarily divulges all sorts of intimate, personal, and undocumented information about himself to a friend. She is doubtless exercising control, in a paradigm sense of the term, over personal information about herself as well as over (cognitive) access to herself. But we would not and should not say that in doing so she is preserving or protecting her privacy. On the contrary, she is voluntarily relinquishing much of her privacy. People can and do choose to give up privacy for many reasons. An adequate conception of privacy must allow for this fact. Control definitions do not. (Parent 1983, p. 273)

Parent's idea is that when you voluntarily divulge personal information to a friend, you clearly diminish your privacy with respect to the friend. But control theorists cannot explain this, Parent says, because the person who voluntarily divulges personal information to his friend is in control. However, if control is interpreted as Negative Control, then the person straightforwardly loses control. According to the definition of Negative Control, A's privacy is diminished if A voluntarily gives someone else access to the relevant information *f*. This is exactly what is at stake in the counterexample. The person in Parent's counterexample voluntarily gives his friend access to all sorts of personal information, so according to the Negative Control Account, the person's privacy is diminished.

The second counterexample that follows the A-strategy comes from Leonhard Menges:¹⁷

Now, consider—as a third voluntary divulgence case—a person who has complete control and exercises it by revealing intimate facts to the public. To be realistic, take Peter Railton's admirable Dewey Lecture (2015). In the lecture, Railton presents a series of moments from his personal life that constitute “a transition from insider to outsider, or back” (2015, p. 2). The final transition is constituted by his giving this very talk and then allowing others to upload the manuscript. That's because he talks openly about his depression and, in particular, his “fear of social embarrassment and humiliation” (2015, p. 13). He says: “I now give to all of you my experience, as story, a tale, an example, you might tell others, or yourself, in order to open a non-threatening conversation

¹⁶ These counterexamples are instances of the voluntarily divulgence cases.

¹⁷ Note that Menges does not take this to be a counterexample to the CT. As we shall see in a later section, Menges develops a new version of the CT that he believes to be immune to counterexamples like Parent's.

with yourself or others about what seeking help can do” (Railton 2015, p. 15).
(Menges 2020)

Again, if control is interpreted as Negative Control, then Peter Railton straightforwardly loses control. According to the definition of Negative Control, A’s privacy is diminished if A voluntarily gives someone else access to the relevant information *f*. This is exactly what is at stake in the counterexample. Peter Railton voluntarily gives the audience access to personal information, so according to the Negative Control Account Peter Railton’s privacy is diminished.

Let us now turn to the counterexamples that follow the B-strategy. We begin with three counterexamples that turn on an interpretation of control as Positive Control. The first one comes from Daniel Farber:

To begin with, while voluntariness is an important aspect of privacy, the concept of control requires elaboration. Privacy would seem to cover nudity as an aspect of intimacy; the Peeping Tom is a classic invader of privacy. If privacy includes the right to “control” visual access to one’s body, then it should include not only the right to preclude such access but also the right to allow it. Yet, it seems decidedly odd to say that public indecency laws violate a flasher’s right to privacy. If anything, the flasher seems to be invading the privacy of others with an unwanted intimacy. (Farber 1993, pp. 514–515)

Farber’s idea is that control can plausibly be interpreted so that the flasher does not have control if he is not able to show his body to anyone he wants. The type of control at play here is Positive Control, since it has to do with the flasher wanting to give others access, but failing to do so. The fact that the flasher is not able to show his body to others is sufficient for Positive Control to be lost, but not sufficient for Negative Control to be lost. If the control theorist makes clear that control on the CT does not mean Positive Control, then Farber’s counterexample does not show that the CT is too broad, since then it does not follow that control is lost just because the flasher is not able to show his body to anyone he wants. Farber’s example has bite only if control is interpreted as Positive Control.

The second counterexample that turns on an interpretation of control as Positive Control comes from Steve Matthews:

A man might not be able to reveal some private information about himself, even if he wants to. Imagine he suffers from temporary dumbness just as he is about to tell his friends about his love life. In such a case, it doesn’t appear that he suffers from a loss of privacy, even though he seems to lack the capacity to reveal his private information. In this case the man retains privacy but lacks control. (Matthews 2008, p. 141)¹⁸

Matthews’s idea is that control can plausibly be interpreted so that the man does not have control if he is not able to reveal his private information to his friends. The type

¹⁸ According to a footnote in (Matthews 2008, p. 141), Matthews got this example from Daniel Cohen in a personal correspondence.

of control at play here is Positive Control, since it has to do with the man wanting to give others access, but failing to do so. The fact that the man is not able to reveal private information to his friends is sufficient for Positive Control to be lost, but not sufficient for Negative Control to be lost. If the control theorist makes clear that control on the CT does not mean Positive Control, then Matthews's counterexample does not show that the CT is too broad, since then it does not follow that control is lost just because the man is not able to reveal private information to his friends. Matthews's example has bite only if control is interpreted as Positive Control.

The third counterexample that turns on an interpretation of control as Positive Control comes from Jeffrey Reiman:

... it might be objected that I can after all invite someone to watch me perform my excretory functions, and in this sense even the privacy that I have here includes my control over who gets access to me. But to think that this shows that such privacy necessarily includes control, one would have to maintain that if I couldn't invite a witness in to watch (say, because of draconian laws or unfailing taboos against doing so), that would mean that those functions were no longer shielded by privacy—and that sounds quite implausible. (Reiman 1995, pp. 30–31)

Reiman's idea is that control can plausibly be interpreted so that you do not have control if you are not able to make other people watch you perform your excretory functions. The type of control at play here is Positive Control, since it has to do with wanting to give someone access, but failing to do so. The fact that you cannot succeed in making other people watch you perform your excretory functions is sufficient for Positive Control to be lost, but not sufficient for Negative Control to be lost. If the control theorist makes clear that control on the CT does not mean Positive Control, then Reiman's counterexample does not show that the CT is too broad, since then it does not follow that control is lost just because you are not able to make other people watch you perform your excretory functions. Reiman's example has bite only if control is interpreted as Positive Control.

Let us now proceed to three counterexamples that turn on an interpretation of control as Republican Control.¹⁹ The first one also comes from Reiman, and it appears immediately before the quote above. Reiman writes:

If it is said that such prohibition [of performing the excretory functions in public] doesn't take away your ability to display such functions [the excretory functions], it only ups the cost of doing so, then it will follow that no one has any privacy in his home since crooks can break in even though it is prohibited. (Ibid.)

Reiman's idea is that control can plausibly be interpreted so that you do not have control if crooks are able break into your house. The type of control at play here

¹⁹ For many more counterexamples that turn on an interpretation of control as Republican Control, see the ones discussed in (Davis 2009, pp. 456–457), and the ones in (Rickless 2007, pp. 782–786). All of these counterexamples are instances of the threatened loss cases.

is Republican Control, since it has to do with the crooks being able to access if they want to. The fact that crooks are able to break in is sufficient for Republican Control to be lost, but not sufficient for Negative Control to be lost. If the control theorist makes clear that control on the CT does not mean Republican Control, then Reiman's counterexample does not show that the CT is too broad, since then it does not follow that control is lost just because crooks are able to break in. Reiman's example has bite only if control is interpreted as Republican Control.

The second counterexample that turns on an interpretation of control as Republican Control comes from Judith Jarvis 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. (Thomson 1975, p. 304)

Thomson's idea is that control can plausibly be interpreted so that you do not have control if your neighbor invents an X-ray device which enables him to look through walls. The type of control at play here is Republican Control, since it has to do with the neighbor being able to access if she wants to. The fact that the neighbor is able to look through the wall if she wants to is sufficient for Republican Control to be lost, but not sufficient for Negative Control to be lost. If the control theorist makes clear that control on the CT does not mean Republican Control, then Thomson's counterexample does not show that the CT is too broad, since then it does not follow that control is lost just because your neighbor invents the X-ray device. Thomson's example has bite only if control is interpreted as Republican Control.

The third counterexample that turns on an interpretation of control as Republican Control comes from Kevin Macnish:

Imagine that I have returned to the coffee shop after a 30 minute interval to find my diary on the table. It is unopened. 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. I feel an enormous sense of relief, thank her and leave with my dignity intact. In this case, I do not think that my privacy has been lessened. When I see my diary in another's possession, I fear that my privacy has been violated, and indeed it might have been. However, as long as the diary is not actually opened and read no reduction in privacy has occurred. Note that this is true even though the diary was not under my control for 30 minutes. (Macnish 2018, pp. 421–422)

Macnish's idea is that control can plausibly be interpreted so that you do not have control if you forget your diary on the table in a coffee shop. The type of control at play here is Republican Control, since it has to do with the stranger being able

to access if she wants to. The fact that the stranger is able to look in the diary is sufficient for Republican Control to be lost, but not sufficient for Negative Control to be lost. If the control theorist makes clear that control on the CT does not mean Republican Control, then Macnish's counterexample does not show that the CT is too broad, since then it does not follow that control is lost just because the stranger *could* read the diary. Macnish's example has bite only if control is interpreted as Republican Control.

I have now given two examples of how the Negative Control Account averts counterexamples that follow the A-strategy, and six examples of how it averts counterexamples that follow the B-strategy. I believe that these points generalize to *any* attempt to construct a counterexample to the CT that follows either the A-strategy or the B-strategy, respectively. If the critics of the CT can give a counterexample that follows either of these strategies, and presupposes a notion of control that is Negative Control, then they have provided a genuine counterexample to the CT. Unfortunately, as we shall see in the next section, interpreting control as Negative Control collapses the CT into the Access Theory (AT).

How Negative Control Collapses the Control Theory into the Access Theory

It should be clear by now that the control theorist can avert the classic counterexamples if she simply points out that control should be interpreted as Negative Control. Although this point holds regardless, it is interesting to consider the prospects of interpreting control as Negative Control. Interpreting control as Negative Control solves many problems for the control theorist, but if doing so introduces new problems, then at least this is a relevant consideration for the control theorist. I believe that defining privacy in terms of Negative Control collapses the CT into the AT.²⁰ Historically, the AT has been the main rival to the CT (Macnish 2018). So, if interpreting control as Negative Control collapses the CT into the AT, then this is very worrying for the control theorist. The control theorist must either accept this collapse, or come up with an alternative interpretation of control that avoids the collapse.

The AT comes in many different versions, but common to all of them is the idea that actual access to the personal information in question is both necessary and sufficient for an individual's privacy to be diminished:

²⁰ Lundgren makes a structurally similar move when he argues that the problems of the CT can only be averted by giving up the concept of control in favor of the concept of limited access (Lundgren 2020, p. 172).

The Access Theory (AT)

An individual A has informational privacy relative to another individual B and to a personal fact f about A if and only if B does not actually access f .²¹

Just like the CT, the AT comes in many different variations. A common commitment among access theorists, however, seems to be that others must actually access the information in question in order for privacy to be diminished (Macnish 2018, p. 421). A crucial motivation behind the AT is the idea that control does no work in determining whether someone has privacy or not. All that matters, according to the AT, is whether someone actually accesses f . It has recently been argued that the access in question is best understood as an actual epistemic access, and that the degree to which A's privacy is diminished depends *inter alia* on how strong the epistemic relation is between B and f (Blaauw 2013; Matheson 2007; Kappel 2013; Fallis 2013). Nothing of importance hangs on whether this specification of the AT is true, but for present purposes, it is helpful to think of the access in question as an actual epistemic access.

Here is how the CT collapses into the AT, if control is interpreted as Negative Control: According to the definition of Negative Control, it is a necessary condition for A's privacy to be diminished that someone else actually accesses f . To see this, recall that there are two ways to diminish Negative Control. The first way of diminishing privacy occurs if (i) and (ii) are not satisfied. This involves someone else attempting to access f , and succeeding because A cannot prevent it. In that case, f is accessed. An example of this would be if a hacker gains access to A's online diary, despite A's best efforts to keep the hacker from accessing. The second way of diminishing privacy occurs if (iii) is not satisfied. This involves A voluntarily letting someone else access f . In that case too, f is accessed. An example of this would be if A voluntarily sends a copy of the diary to the hacker. So, either way, if A loses Negative Control, then someone has accessed f . This makes the Negative Control Account completely coextensional with the AT. Therefore, an access theorist can insist that what drives our intuitions when we think that A's privacy is diminished is the fact that someone accesses f , rather than the fact that A loses control over f .

By saying that the Negative Control Account becomes coextensional with the AT, I mean the following: In any given case, if the Negative Control Account gives the verdict that A's privacy is diminished, the AT also gives this verdict. When I say that this collapses the Negative Control Account into the AT, I do not mean that the Negative Control Account gives the same verdicts as the AT *for the same reasons*. Following Menges, I mean only that the Negative Control Account and the AT are coextensional in the way described above (Menges 2020, p. 3), and that this gives the access theorist room to insist that in any given case, what explains A's

²¹ Variations of the AT can be found in (Thomson 1975), (Gavison 1980), (Bok 1989), (Allen 1988), (van den Haag 1971), (Reiman 1995), (Macnish 2018), (Lundgren 2020), and others.

diminishment of privacy is the fact that someone else accesses f , and not the fact that A loses Negative Control over f .²²

Now, a control theorist might insist that CT is not coextensional with the AT, because there are cases where privacy is diminished even though no one accesses any personal information. But it is difficult to see how such a case could be constructed without making the exact same mistake as critics of the CT have made; namely interpreting control as Republican Control. Any case where f is not actually accessed, but where A lacks some sort of control over f , seems to be akin to the threatened loss cases provided by Thomson, Macnish, etc., where A does not have Republican Control. Thus, this type of reply is not available to the control theorist if she wants to avert the counterexamples that follow the B-strategy in the way that I have suggested in this paper. In other words, the control theorist cannot define Negative Control in a way that avoids the collapse if they also want to maintain that privacy is not diminished just because A does not have Republican Control.

The control theorist might insist instead that there are cases where privacy is diminished even though no one accesses any personal information, and where this verdict does not rely on the republican interpretation of control. For instance, the control theorist might point to something like Jakob Mainz and Rasmus Uhrenfeldt's recent Wiretapping case:

Wiretapping

Smith and Jones are neighbors. Unbeknownst to Jones, Smith wiretaps Jones' telephone, using a fancy device which allows Smith to listen in on Jones' conversations without violating Jones' property rights. As it happens, Jones is on vacation for several months, and does therefore not use the telephone in that time period. (Mainz and Uhrenfeldt 2020)

Wiretapping is meant to elicit the intuition that a violation of the right to privacy can occur, even if no one gets access to personal information. Smith does not access personal information about Jones, because Jones happens not to use the telephone. Nevertheless, it might seem as if Smith violates Jones's right to privacy (*ibid.*). The control theorist might point to something like Wiretapping to explain why the CT does not collapse into the AT if control is interpreted as Negative Control. At least on one reading of Negative Control, Jones's privacy is diminished, because conditions (i) and (ii) are not satisfied. Smith attempts to get access, but the reason why he fails is not because of Jones's intentional actions directed at preventing others from accessing. On this reading of Negative Control, it is not coextensional with the AT. However, if an implication of interpreting control as Negative Control is that Jones's

²² Menges claims that Lundgren's argument is meant to show that the CT collapses into the AT. However, Lundgren never calls it a collapse. Supposedly, we can have a weak and a strong sense of collapsing. According to the weak one, a theory collapses into another theory if the first theory gives the same verdict as the second one. According to the strong one, a theory collapses into another theory if the first theory gives the same verdict as the second one, *for the same reasons*. Like Menges, I follow the weak sense of collapsing, when I say that the Negative Control Account collapses into the AT.

privacy is diminished in Wiretapping, then this counts *against* the Negative Control Account. Wiretapping is meant to show that a *violation* of the right to privacy can occur even if no one accesses personal information. It does *not* show that a *diminishment* of privacy can occur even if no one accesses personal information. It is not clear at all that Jones's privacy is diminished in Wiretapping.²³ So, if the control theorist insists that the Negative Control Account does not collapse into the AT because the former implies that there is no diminishment of privacy in Wiretapping, then so much the worse for the control theorist.

The access theorist, on the other hand, can straightforwardly insist that Jones's privacy is not diminished because Smith does not get access to any personal information about Jones. In fact, it is not clear either that a violation of the right to privacy occurs in Wiretapping. It seems more intuitive to say that what happens in Wiretapping is an *attempt* to violate Jones's right to privacy. In order for this attempt to succeed, Smith would have needed to actually access Jones's personal information, which he did not. In this counterfactual case, it would also be the case that Jones's privacy is diminished. But then, the access theorist could plausibly reply that this loss of privacy occurs exactly because Smith accesses Jones's personal information. It therefore seems that the control theorist needs to accept that the CT collapses into the AT, if she interprets control as Negative Control.²⁴ In the next section, I will present an objection to my argument.

The Source Control Objection

In this section, I will discuss an objection to my argument. According to this objection, the control theorist can avoid the collapse into the Access Theory, if she interprets control as 'source control' instead of Negative Control. As I will show in this section, the source control account should be rejected for independent reasons. Thus, even if this objection is true, the control theorist should not interpret control as source control.

As mentioned in the introduction, Leonhard Menges has recently defended a version of the CT that he calls the 'source control account of privacy'. The account is novel, and suggests a promising alternative answer to the question of how the control theorist should interpret 'control'. Menges argues—although he does not use this terminology—that his account can both avert the

²³ Except perhaps with regards to the information that Jones does not use the phone. But note that Smith *does* access this information.

²⁴ An anonymous reviewer suggested to me that the point about collapse is not very interesting because Negative Control is defined in a way that makes it obvious that losing Negative Control entails that someone has access. However, I believe that this is a feature, not a bug. The point is exactly that if the control theorist wants to avoid all the classic counterexamples against the CT, then she needs to define control in a way that entails that a loss of control entails access. To see this, consider the implications of removing the parts of the definition that makes it obvious that a loss of control entails that someone has access. What you will find is that removing these parts of the definition simply reactivates some of the classic objections against the CT.

counterexamples to the CT that follow the A-strategy and those that follow the B-strategy. He also argues that his account does not collapse into the AT. If all of this is correct, then the control theorist can follow Menges's account instead of the Negative Control Account, and thus avoid the collapse into the AT. However, I believe that there is reason to think that Menges's account does in fact not avert all counterexamples that follow the A-strategy.

According to Menges, control theorists should interpret control as what he calls 'source control'. This notion of control is inspired by the classic Frankfurt-cases known from the literature on free will, such as the following:

Jones has resolved to shoot Smith. Black has learned of Jones's plan and wants Jones to shoot Smith. But Black would prefer that Jones shoot Smith on his own. However, concerned that Jones might waver in his resolve to shoot Smith, Black secretly arranges things so that, if Jones should show any sign at all that he will not shoot Smith (something Black has the resources to detect), Black will be able to manipulate Jones in such a way that Jones will shoot Smith. As things transpire, Jones follows through with his plans and shoots Smith for his own reasons. No one else in any way threatened or coerced Jones, offered Jones a bribe, or even suggested that he shoot Smith. Jones shot Smith under his own steam. Black never intervened. (McKenna and Coates 2020, sect. 3.2).

Menges explains that while Jones could not have avoided killing Smith, Jones still exercises an important kind of control when he decides to shoot Smith without any intervention. We can, as Menges writes, '... have an important kind of control over what we do without having effective choice over whether or not we do it' (Menges 2020, p. 8). The type of control that Jones exercises is what Menges calls source control. If one has this kind of control, then one is the right kind of source of one's actions. Menges leaves it unsatisfactorily unclear what exactly source control is, and he is aware of that (Menges 2020, p. 9). Nevertheless, control theorists should interpret control as source control, Menges says:

My main proposal is that privacy theorists can and should spell out privacy in terms of source control. According to the resulting source control account of privacy, an agent has privacy with regard to a certain piece of information just in case the person is the right kind of source of the relevant information flow if the information flows at all. In other words: 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. (Menges 2020, p. 9)

Menges goes through a series of cases in order to show that his source control account generates the intuitively correct results in all of these cases. I will not go through all of these cases here, but I will note that I agree with Menges that the source control account does avert nicely all the counterexamples that follow the B-strategy. However, Menges also claims that the source control view averts the counterexamples that follow the A-strategy. He argues that the privacy of the person in Parent's counterexample is not diminished, as long as he is the right kind

of source of sharing the information. That is, Menges bites Parent's bullet and says that the privacy of the person in Parent's counterexample is not diminished. He thinks that the person's privacy is not *diminished*, but that he rather *includes* the friend and possibly others in his private realm (Menges 2020, p. 6).²⁵

Like most contemporary privacy scholars, I find this result very counterintuitive in itself. It seems strange that one's privacy is not diminished when one divulges all sorts of personal information to a friend. But I also think that this verdict has counterintuitive implications. For instance, it implies that even if the person voluntarily divulges *all* personal facts about himself to *every living person on earth*, then—as long as he is the right kind of source of sharing the information—his privacy is not diminished even a tiny bit. That is, the person has full privacy with respect to everyone, even though everyone knows everything about him. This seems like a very counterintuitive result. To illustrate, consider the following thought experiment:

Moving Day

Every citizen of Private Ville lives in regular houses made of bricks. Every citizen of Private Ville is being wiretapped against his or her will by someone from outside of Private Ville. One day, every citizen of Private Ville chooses to move to houses that are made of fully transparent glass. Everyone that walks by such a house can see everything that happens inside the house. And, because the walls are made of thin glass, everyone outside the house can also hear every little sound from inside the house. No one is wiretapping the citizens of Private Ville in the new houses. But the people who were doing the wiretapping are now standing outside the glass houses, watching and listening to what citizens of Private Ville do inside their houses. The citizens of Private Ville are fully aware of this.²⁶

On the assumption that every citizen of Private Ville exercises source control when they choose to live in such a house, and when they choose to say and do things within the house, it follows from the source control account that by moving into the glass houses the citizens of Private Ville are performing a privacy *enhancing* action.²⁷ They go from not having privacy with regard to information about what happens inside their houses—because they did not have source control when they were wiretapped against their will—to having *full* privacy with regard to this information—because they now exercise source control. This seems very strange. The information still flows to the outsiders of Private Ville as before, but now it *also*

²⁵ See (Inness 1992, p. 46) for a similar reply to Parent.

²⁶ Moving Day is inspired by an example from (Floridi 2006, p. 110).

²⁷ Menges argues one way to think about source control in relation to privacy is that a person is the right kind of source of the information flow if the person has a first-order desire that the information flows to others, and a second-order desire that she has the first-order desire (Menges 2020, p. 9). On this version of source control, I would need to say that the citizens in Private Ville have both first-order and second-order desires that the information about what they do inside the glass houses flows to the people outside the houses. However, Menges also says that this is not the view he argues for (ibid.).

flows to everyone else who happens to walk by. And yet, the source control account implies that the citizens of Private Ville now have *more* privacy than before.

Note that even if this is not a completely counterintuitive result, at least it seems that Menges's source control account does not handle counterexamples that follow the A-strategy nearly as straightforwardly as the Negative Control Account does. The source control account, and the Negative Control Account, seem to handle counterexamples that follow the B-strategy equally well. But the Negative Control Account handles counterexamples that follow the A-strategy much more straightforwardly than the source control account does. Recall that on the Negative Control Account, the privacy of the person in Parent's counterexample is diminished because part iii) of the definition of Negative Control is not satisfied. Thus, the Negative Control Account can straightforwardly handle voluntary divulgence cases like Moving Day. So, all things being equal, the Negative Control Account seems more promising than the source control account when it comes to handling the counterexamples to the CT that follow the A-strategy. It is therefore all the more problematic for the control theorist that the Negative Control Account collapses into the AT.

Now, control theorists have often replied to Parent that on the CT, the privacy of the person is in fact diminished, because the person loses control over whether the friend will distribute the personal information to others (Gavison 1980, p. 427; Matheson 2007, p. 255; Lundgren 2020, pp. 168–169). While this reply is intuitively appealing, it is not available to Menges if he wants to remain consistent. The reason is that this reply to Parent claims that the privacy of the person who voluntarily divulges personal information to a friend *is* diminished. But, as we have seen, Menges explicitly denies this. So, Menges cannot fall back on this reply if he wants to remain consistent.

It may be true that the control theorist can avoid the collapse into the AT, if control is interpreted as source control rather than Negative Control. But, the control theorist should not interpret control as source control regardless, because the source control account cannot handle the counterexamples that follow the A-strategy. This leaves the control theorist with the Negative Control Account which—as we have seen—collapses into the AT.

Concluding Remarks

In this paper, I have defended the AT. I have done so indirectly by developing a new version of the rival CT that is immune to all the classic objections, and showing how this version collapses into the AT. The novel distinction between three types of control, Negative Control, Positive Control, and Republican Control respectively, allows the control theorist to avert both the counterexamples to the CT that follow the A-strategy, and those that follow the B-strategy. If the control theorist points out that control should be interpreted as Negative Control, then all the counterexamples lose their bite. This result itself helps clear up the messy and extensive literature on how best to define privacy. I believe that I have identified a way to save the CT from the most worrying counterexamples against it. This is not a victory for the control theorist, however, given the solution implies that the CT collapses into the AT.

I have discussed a recent version of the CT—the source control account—that does not collapse into the AT. This version does not avert the counterexamples that follow the A-strategy, though. Moving forward, this leaves three options available to the control theorist: The first option is to admit defeat because the Negative Control Account collapses into the AT. The second option is to follow the source control account and look for more plausible ways to handle the counterexamples that follow the A-strategy. The third option is to look for a third version of the CT that handles all the counterexamples and does not collapse into the AT. Regardless of which direction the discussion goes, I believe that progress is made.

Declarations

Conflict of interest There are no conflicts of interests to declare.

References

- Allen, Anita. 1988. *Uneasy access: Privacy for women in a free society*. Rowman & Littlefield Publishers.
- Altman, Irwin. 1976. Privacy: A conceptual analysis. *Environment and Behavior* 8 (1): 141–141.
- Benzanson, Randall. 1991. Privacy, personality, and social norms. *The Case Western Reserve Law Review*. 41 (3): 681–687.
- Berlin, Isaiah. 1969. *Two concepts of liberty*. Clarendon Press.
- Birnhack, Michael. 2019. A process-based approach to informational privacy and the case of big medical data. *Theoretical Inquiries in Law* 20: 257–290.
- Blaauw, Martijn. 2013. The epistemic account of privacy. *Episteme* 10 (2): 167–177.
- Bok, Sissela. 1989. *Secrets: On the ethics of concealment and revelation*. New York: Vintage Books.
- Calo, M. Ryan. 2011. The boundaries of privacy harm. *Indiana Law Journal* 86 (3): 1131–1162.
- Davis, Steven. 2009. Is there a right to privacy? *Pacific Philosophical Quarterly* 90: 450–475.
- Fallis, Don. 2013. Privacy and lack of knowledge. *Episteme* 10 (2): 153–166.
- Falls-Corbitt, Margaret, and Michael McLain. 1992. God and Privacy. *Faith and Philosophy* 9 (3): 369–386.
- Farber, Daniel. 1993. Book Review: Privacy, Intimacy, and Isolation. By Julie C. Inness. *Constitutional Commentary* 198: 510–519.
- Floridi, Luciano. 2006. Four challenges for a theory of informational privacy. *Ethics and Information Technology* 8: 109–119.
- Frey, R. 2000. Privacy, control, and talk of rights. *Social Philosophy and Policy* 17 (2): 45–67.
- Fried, Charles. 1968. Privacy. *Yale Law Journal* 77 (3): 475–493.
- Froomkin, Michael. 2000. The death of privacy? *The Stanford Law Review* 52: 1461–1544.
- Gavison, Ruth. 1980. Privacy and the limits of law. *The Yale Law Journal* 89 (3): 421–471.
- Goldberg, Ian, Austin Hill, and Adam Shostack. 2001. Trust, ethics and privacy. *Boston University Law Review* 81 (2): 407–422.
- Gross, Hyman. 1971. Privacy and autonomy. In *Nomos XIII: Privacy*, pp 169–181.
- Hoye, J. M., and Jeffrey Monaghan. 2018. Surveillance, freedom and the republic. *European Journal of Political Theory* 17(3): 343–363.
- Inness, Julie. 1992. *Privacy, intimacy, and isolation* Oxford University Press.
- Kappel, Klemens. 2013. Epistemological dimensions of informational privacy. *Episteme* 10 (2): 179–192.
- Lundgren, Björn. 2020. A dilemma for privacy as control. *The Journal of Ethics* 24: 165–175.
- Macnish, Kevin. 2018. Government surveillance and why defining privacy matters in a post-Snowden world. *Journal of Applied Philosophy* 35(2): 417–432.
- Mainz, Jakob, and Rasmus Uhrenfeldt. 2020. Too much info: Data surveillance and reasons to favor the control account of the right to privacy. *Res Publica* 27 (2): 287–302.

- Matheson, David. 2007. Unknowableness and informational privacy. *Journal of Philosophical Research* 32: 251–267.
- Matthews, Steve. 2008. Privacy, separation, and control. *The Monist* 91 (1): 130–150.
- Menges, Leonhard. 2020. A defense of privacy as control. *The Journal of Ethics*. Online first.
- Miller, Seuman, and John Weckert. 2000. Privacy, the workplace and the internet. *Journal of Business Ethics* 28: 255–265.
- Moore, Adam D. 2003. Privacy: Its meaning and value. *American Philosophical Quarterly* 40 (3): 215–227.
- Moore, Adam. 2010. *Privacy rights: Moral and legal foundations*. Pennsylvania State University Press.
- Newell, Bryce Clayton. 2018. Privacy as antipower. In pursuit of non-domination. *European Data Protection Law Review* 4 (1): 12–16.
- Parent, William. 1983. Privacy, morality and the law. *Philosophy & Public Affairs* 12 (4): 269–288.
- Parker, Richard. 1974. A definition of privacy. *Rutgers Law Review* 27: 275–296.
- Pettit, Philip. 1999. *Republicanism: A theory of freedom and government*. Oxford University Press.
- Rachels, James. 1975. Why privacy is important. *Philosophy & Public Affairs* 4 (4): 323–333.
- Reiman, Jeffrey H. 1995. Driving to the panopticon: A philosophical exploration of the risks to privacy posed by the highway technology of the future. *Privacy*. 11: 159–176.
- Rickless, Samuel. 2007. The right to privacy unveiled. *San Diego Law Review* 44 (4): 773.
- Roberts, A. 2014. A republican account of the value of privacy. *European Journal of Political Theory* 14 (3): 320–344.
- Roessler, Beate. 2005. *The Value of Privacy*. Polity.
- Solove, Daniel J. 2002. Conceptualizing privacy. *California Law Review* 90 (4): 1087–1156.
- Thomson, Judith Jarvis. 1975. The right to privacy. *Philosophy & Public Affairs* 4 (4): 295–314.
- van den Haag, Ernst. 1971. On Privacy. In *Privacy: Nomos XIII*, 149–168.
- van der Sloot, Bart. 2018. A new approach to the right to privacy, or how the European court of human rights embraced the non-domination principle. *Computer Law and Security Review* 34 (3): 539–549.
- Warren, Samuel, and Louis Brandeis. 1890. The right to privacy. *Harvard Law Review* 4 (5): 193–220.
- Westin, Alan F. 1970. *Privacy and Freedom*. Bodley Head.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

ARTICLE



But anyone can mix their labor: a reply to Cheneval

Jakob Thrane Mainz 

Center for Applied Philosophy, Aalborg University, Aalborg, Denmark

ABSTRACT

Francis Cheneval has recently argued that people have property rights over personal data about themselves. Until now, the discussion on data ownership has primarily been a discussion among legal theorists and economists. Cheneval contribution to the discussion is a very welcome input from academic philosophy. Cheneval attempts to reach his conclusion through two distinct strategies. One strategy is to reach the conclusion through a Lockean inspired libertarian rights-based theory of property. The second strategy is to reach his conclusion through a Rawlsian account of distributive justice. According to Cheneval, his conclusion can be reached both ways. In this reply, I will focus exclusively on Cheneval argument that people have Lockean inspired libertarian property rights over personal data. I will offer an objection, which – if correct – demonstrates how Cheneval Lockean argument runs into a dilemma.

KEYWORDS Data ownership; Locke; Property rights; Personal data

Francis Cheneval has recently argued that people have property rights over personal data about themselves. Until now, the discussion on data ownership has primarily been a discussion among legal theorists and economists. Cheneval's contribution to the discussion is a very welcome input from academic philosophy. Cheneval attempts to reach his conclusion through two distinct strategies. One strategy is to reach the conclusion through a Lockean inspired libertarian rights-based theory of property. The second strategy is to reach his conclusion through a Rawlsian account of distributive justice. According to Cheneval, his conclusion can be reached both ways. In this reply, I will focus exclusively on Cheneval's argument that people have Lockean inspired libertarian property rights over personal data. I will offer an objection, which – if correct – demonstrates how Cheneval's Lockean argument runs into a dilemma.

In section A, I outline Cheneval's Lockean argument. In section B, I present my objection to Cheneval's Lockean argument in the form of a dilemma. In section C, I discuss the prospect of solving the dilemma. Finally, in section D, I conclude.

CONTACT Jakob Thrane Mainz  jtm@learning.aau.dk

© 2020 Informa UK Limited, trading as Taylor & Francis Group

Cheneval's argument

Cheneval argues that people have property rights¹ over personal data about themselves, and that these property rights should be acknowledged by law.² Granting legal property rights over personal data would give people far-reaching control over what happens to certain information about them (Cheneval, 2018, p. 3). An ownership regime like this would make it possible for people to hold a data bank, and buy and sell personal data on a free market. Such a data market might, according to Cheneval, correct certain misallocations in the 'data economy', and it could contribute to the financing of people's pensions (Cheneval, 2018, 2).

However, Cheneval does not think that people have property rights over *all* personal data, since this view would have strange implications, according to Cheneval (Ibid.). He writes:

For instance, Jones sees Smith eating in a restaurant at a certain point in time. It would be strange to argue that Smith holds privacy and property rights over the information Jones has stored about him in her brain after seeing him in the restaurant and that Smith alone can determine what Jones is allowed to do with that information. The meaningful object of privacy and property rights, especially if they are to be cast in law, is not personal information as such, but the way in which it is obtained, registered, certified, re-used, aggregated, made accessible and so forth. (Ibid, 3)

In order to avoid the strange implication, he delimits the scope of the property rights to include only what the EU includes in their definition of personal data: Personal data means data relating to an identified or identifiable natural person (data subject). This can for example, be a name, an identification number, location data, an online identifier, or a physical, physiological, mental, economic, cultural or social identity of that natural person (Union & regulation, 2016/67, article 4,1). It is not clear to me how this delimitation avoids the strange implication. It could plausibly be argued that Jones does indeed get access to e.g., location data about Smith, when Jones sees Smith in the restaurant. So, either the example of Smith and Jones is not a good example, or Cheneval's delimitation of the scope of what counts as personal data is problematic. Either way, it is not completely clear what the last sentence in the quote above means. What does it mean to own 'the way in which it is obtained, registered, certified, re-used, aggregated, made accessible and so forth'? On a charitable reading of this, it seems that what Cheneval has in mind is that people can come to own personal data when they are 'processing' the personal data in at least the following ways: 1) By obtaining the data, 2) by registering the data, 3) by certifying the data, 4) by re-using the data, 5) by aggregating the data, or 6) by making the data accessible.

Who can come to own personal data these ways? Cheneval thinks – uncontroversially – that people can come to own personal data about themselves. But, if a ‘counter-party’, as Cheneval calls it, processes the data in question, then this counter-party can obtain partial ownership of the data as well. Nevertheless, the person whom the data is about necessarily one of the owners of the data in question. Cheneval writes:

... for now it suffices to say that most schemes of legitimate private property of personal data will be arrangements of co-ownership with different bundles of rights in the hands of interactive co-owners. However, the point is that the natural person at the beginning of the value chain, however small her contribution, is necessarily one of the right holders of property rights of her personal data. (Cheneval, 2018, p. 4)

We can derive two important claims from this quote. The first claim is that it is possible to have co-ownership in personal data. The second claim is that the person who the data is about is necessarily one of the owners of the data. In order to support these claims, he appeals to the two Lockean ideas that i) ‘mixing labor’ with the data (by processing the data in the ways described in 1–6) generates *prima facie* ownership over the data, and ii) that people own themselves and thus also own data about themselves:

Building on this argument [a Lockean inspired libertarian argument] the property claim to personal data follows from the intuition that persons are the original owners of their personhood, bodies and minds, and hence of information that constitutes their personal identity and/or that is generated by their private data registration activities. Data that persons generate by registering their name and address, by engaging in transactions that leave decipherable traces, etc. are therefore *prima facie* the personal property of those persons in question. If information is digitally processed by the person’s own activity or by her participation in such activity, and if it specifically refers to a person as an individual, it is to be in the ownership and control of the person in question. If registered information on individual persons and personal activities is used in the activities of others, investing labour and capital in an economic endeavour, persons first of all ought to have a say what can be done with their personal data and they ought to have a partial claim to the benefits stemming from the economic activities that use their personal data as a resource. (Ibid)

The idea is that if you mix your labor with some personal data, be it data about yourself or someone else, then you get at least *prima facie* partial ownership over these data. And, since people are the ‘original owners of their personhood, bodies and minds’, they are also the owners of ‘information that constitutes their personal identity’. It is clear that these ideas have a Lockean flavor, but for the sake of overview, let us see how exactly Cheneval’s thinks his idea of data ownership can be derived from a Lockean account of property. Here is a standardized version of Locke’s original argument:

The Lockean Argument³

Premise 1: If persons are the original owners of their respective personhoods, bodies and minds, then mixing their labor with something unowned generates property rights over the thing in question (provided that a certain proviso is satisfied).

Premise 2: Persons are the original owners of their respective personhoods, bodies and minds.

Conclusion: Mixing a person's labor with something unowned generates property rights over the thing in question (provided that a certain proviso is satisfied).

Cheneval's idea is that data ownership can be justified by applying Premise 1 and/or Premise 2 of the Lockean Argument to personal data, as we saw in the quote above. Applying Premise 1 to personal data means to refer to the intuition that certain types of data processing (at least the ones in 1–6) constitute mixing of labor with the data in question. Applying Premise 2 to personal data means to invoke what has often been called the Self-Ownership Thesis (SOT). The SOT is the antecedent in Premise 1: Persons are the original owners of their respective personhoods, bodies and minds. Cheneval's idea seems to be that data *about* the self is part of what constitutes this self, and that the SOT thus applies to personal data too.⁴ Cheneval defends Premise 2 in the Lockean Argument against some of the prominent objections against the SOT as such. But he does not consider *any* objections against the view that the SOT applies to personal data. Neither does he consider *any* objections to the view that Premise 1 applies to personal data. If we stay true to Premise 1, it seems that there is a gap between *prima facie* property rights and full-blown property rights in Cheneval's argument. This gap can presumably be closed by adding a standard Lockean proviso, which says that 'enough-and-as-good' must be left for others.⁵ In relation to personal data, this proviso seems easily satisfied. There will certainly be 'enough-and-as-good' personal data about others left for them to obtain ownership over, especially when we consider that information is generally a non-rivalrous. This means that – contrary to tangible things like shoes – Smith can use his personal information while Jones also uses it. We have now seen how Cheneval thinks that the Lockean Argument applies to personal data. In the next section, I will present my objection to Cheneval's argument.

The objection

My main objection is this: Applying the Lockean Argument to personal data in the way Cheneval does leads to a dilemma. The dilemma consists in choosing between applying Premise 1 in the Lockean Argument to personal data, and applying Premise 2 in the Lockean Argument to personal data. Let us first consider an implication of accepting that Premise 1 applies to personal data.

It follows from Cheneval's argument that *anyone* can mix their labor with data about you, before you do it, and thus obtain at least partial⁶ property rights over these data about you. This has wildly counterintuitive implications. Let us see how this follows from Cheneval's argument. Cheneval writes the following conditional: 'If information is digitally processed by the person's own activity or by her participation in such activity, and if it specifically refers to a person as an individual, it is to be in the ownership and control of the person in question.' (Ibid, 8). In this quote, Cheneval claims that the fact that someone, call him Smith, *participates* in the processing of personal data about Smith, is a sufficient condition for Smith obtaining property rights over the personal data in question. But, as we saw earlier, Cheneval does not rule out that Jones, who also participates in the processing of Smith's data, can obtain at least partial ownership over the personal data about Smith too. Jones can get partial ownership over personal data about Smith, if Jones mixes his labor with these data. And, if Jones makes money off of these data, then Smith has a claim to at least some of the money. The question is now: What happens if *Smith* does not mix any labor with the data in question, while *Jones* does? In that case, it seems that Cheneval – straightforwardly applying Premise 1 to personal data – is committed to the view that Jones is now the sole owner of the personal data about Smith. This seems very counterintuitive to me. Let us consider an example:

Restaurant

Smith is a very famous actor. He is having dinner with his friend Jones at a restaurant. Jones is a freelance journalist, and unbeknownst to Smith, Jones is covering Smith's everyday life. Jones is secretly transcribing the entire conversation on his tablet. Smith reveals to Jones the fact that Smith has terminal cancer. The day after, Jones sells the transcript to the tabloid press.

According to Cheneval's argument, Jones has obtained at least partial ownership over the personal data about Smith by processing the data, and he can now rightfully transfer his ownership to the tabloid press. If Smith has not participated in the processing of the data, then Smith has no property claim to them. But, if Peter, Carl or Allan also start processing the personal data about Smith, then they obtain at least partial ownership over the data. There is no limit to how many people can obtain ownership over Smith's data, since data is a non-rivalrous good. The only way in which Smith can gain at least partial ownership over the personal data is by processing the data himself. And even in that case, Peter, Carl and Allan can still get partial ownership too.⁷

Now, there are at least three replies available to Cheneval, but I think that none of them work. The first reply available to Cheneval is this: As we saw earlier, one of the ways in which one can mix labor with personal data is by making the data 'accessible'. In Restaurant, Smith is indeed making the data

about his disease accessible to Jones by talking about it, so Smith owns the data before Jones transcribes it. Therefore, Smith does in fact have at least partial ownership over the data. Here is why I think this reply does not work: Imagine a slightly altered version of Restaurant. In this version Jones is a doctor, and he sees Smith's symptoms and writes down information about them, before Smith has told Jones anything about the disease. So, Smith has not mixed his labor with the data before Jones has mixed his labor with it. Should we now accept that Jones owns the data, and Smith does not? Clearly not. Cheneval might then reply that even having visible symptoms counts as making the data about the disease accessible, and therefore Smith owns the data before Jones the doctor starts transcribing. However, if having visible symptoms counts as making the data about the disease accessible, then it seems that the idea of 'making accessible' collapses into an idea of self-ownership. If I have a strange disease which causes a set of eyes to grow out of my shoulder,⁸ then I own these eyes because I own myself, not because data about them are made accessible to you.

The second reply available to Cheneval is this: Smith is the original owner of the data, and therefore it does not matter if he mixes his labor or not. After all, recall from one of the quotes earlier that Cheneval thinks that the person who the data is about is 'necessarily one of the right holders of property rights'. If Cheneval is right about this, it seems that Smith does in fact have at least partial ownership over the data. But here is the catch: If Smith has not mixed any labor with the data in question, then it seems very ad hoc to claim that Smith is 'necessarily one of the right holders of property rights'. In order to claim this in a non-ad hoc way, Cheneval needs to invoke Premise 2 in the Lockean Argument and apply it to personal data. If Smith owns the data because they are part of 'him', and he owns himself, then it seems clear that he is the *original* owner of the data. But in that case, if Smith is already the original owner of the data, then mixing labor contributes nothing to Smith's ownership over the data. And neither should Peter, Carl or Allan's mixing of labor with the data entail that they now get partial ownership over it. It is already Smith's data. Consider this analogy: If I am the original owner of my body, then a doctor does not get partial ownership over it, if she mixes her labor with it. If the SOT applies to personal data, then mixing labor plays no role at all.

The third reply available to Cheneval is this: The verdict over Restaurant is correct, but the reason is that Smith has forfeited his property rights over the personal data by talking about the disease in a public restaurant, and the data is therefore up for grabs for anyone who mixes labor with the data. When Peter, Carl and Allan mix their labor with the data, they each obtain partial ownership over the data. This reply, however, presupposes either that Smith was the original owner of the data, or that he owns the data due to mixing labor with it, since he would otherwise not be able to forfeit the right. If Smith

is the original owner of personal data about him, then Cheneval could just have explained data ownership in terms of the SOT alone. If Smith is not the original owner of the data, and if Smith makes *no* contribution to the processing of the data, then it is hard to see how, on Cheneval's own account, Smith can have *any* claim to the data.

In order to avoid the counterintuitive implication exemplified by Restaurant, Cheneval needs to explain data ownership in terms of the SOT alone. But we already saw the problem of this move in the discussion of the second reply above, namely that if he explains data ownership in terms of the SOT alone, then he cannot also explain how people can get partial ownership in personal data merely by mixing labor. Cheneval is thus caught in a dilemma consisting of the following options:

Option 1: Explain data ownership by applying Premise 1 in the Lockean Argument to personal data.

Option 2: Explain data ownership by applying Premise 2 in the Lockean Argument to personal data.

If he chooses Option 1, then his argument has a very counterintuitive implication (exemplified by Restaurant). In order to avoid this implication, he needs to apply Premise 2 to personal data, and thus he must choose Option 2.

If he chooses Option 2, then he cannot also defend data ownership through Premise 1, since an original owner of X does not lose any ownership in X just because someone else mixes labor with X. For this reason, if Cheneval chooses Option 2, then he loses his explanation for how people get partial ownership over personal data about other people, since his explanation consists in applying Premise 1 to personal data.⁹

The most promising way out for Cheneval seems to be Option 2, since Option 1 has the counterintuitive implication exemplified by Restaurant. And in order to avoid this implication, Cheneval needs to choose Option 2. The price of choosing Option 2 is to accept the view that people cannot get ownership over personal data about others simply by mixing labor with the data. This price may be worth paying, though. The next section is devoted to the prospects of Option 2.

The prospects of applying SOT to personal data

In the previous section we saw that the most promising option for Cheneval is Option 2. If people do have property rights in personal data at all, it seems intuitively more plausible that such property rights obtain from the SOT itself and not through the mixing of labor. In other words, it seems more plausible to argue for property rights in personal data by applying

Premise 2 – rather than applying Premise 1 – to personal data. This option avoids the counterintuitive implication spelled out earlier (exemplified by Restaurant), while it still allows for markets in personal data to obtain, since people can still own data and transfer the ownership by engaging in contracts.

Although applying Premise 2 seems *prima facie* more plausible than the applying Premise 1, it is not completely obvious that the SOT applies to personal data. Luciano Floridi has recently defended the view that personal data constitutes the *person*, rather than something *possessed* by the person (Floridi, 2013, p. 243). This is a controversial view which should be developed much further, before we can straightforwardly apply the SOT to personal data. Floridi, for one, does not seem to think that the SOT applies to personal data, even if the person is constituted by these data. In fact, Floridi seems to completely reject the idea that people own personal data (Floridi, 2013, 244). Furthermore, if Cheneval wants to defend data ownership by only applying the SOT to personal data, then the defense is vulnerable to the objections against the SOT in general. Of course, this would also be true if Cheneval wanted to defend data ownership only by applying the idea of mixing labor to personal data, since the idea of mixing labor relies on the SOT. Nonetheless, while Cheneval discusses and rejects several objections to the SOT, he leaves out some of the most hard-hitting ones (see e.g., (Lippert-Rasmussen, 2008)). If Cheneval can make a convincing argument for the claim that the SOT applies to personal data, then it seems that he has a plausible way out of the dilemma. But the price is that mixing labor plays no role when it comes to obtaining property rights over personal data, and therefore Cheneval needs another explanation if he wants to maintain that people can get partial ownership over personal data.

Concluding remarks

In this reply, I have tried to show that Cheneval's Lockean argument in favor of data ownership runs into a dilemma. Cheneval argues that 1) property rights over personal data can be derived through a Lockean process of mixing labor with the data. He also argues that 2) the SOT applies to personal data, and that people thus own data about themselves. I claim that 1) is false, since it has the counterintuitive implication that others can get to own personal data about you, if they mix their labor with it before you do. In order to avoid this implication it must be presupposed that the SOT applies to personal data, and therefore the avoidance of the implication relies on 2). This leaves us with 2), which seems *prima facie* more plausible. But if 2) is true, then Cheneval loses his explanation for how people get partial ownership over personal data about others, since his explanation relies on 1). One way to solve this dilemma is for Cheneval to go with 2) and abandon 1). However,

more work is needed in order to show why we should accept the view that the SOT implies original ownership over personal data.

In all fairness, Cheneval does not rely entirely on the Lockean defense of data ownership. He also claims that the same conclusion can be derived from an argument from Rawlsian distributive justice. If all I have argued in this reply is correct, I have not showed that data ownership as such should be ruled out. I have only showed that Cheneval's Lockean approach is problematic.

The philosophical discussion about data ownership is a very topical and very welcome one, especially considering the emergence of Big Data and data markets in various aspects of our lives. In order to continue the theoretically and practically important discussion of data ownership, I hope that Cheneval will further pursue the issues I have pointed out in this reply.

Notes

1. Throughout this reply, 'property rights' denotes moral property rights, unless explicitly specified otherwise.
2. For recent discussions on data ownership in legal theory and economics, see e.g., (Laudon, 1996); (Samuelson, 2000); (Thouvenin et al., 2017); (Cwik, 2016); (Cohen, 2017).
3. See (Locke, (1690) [1988], ch. 5.)
4. If this is true, then it seems rather ad hoc for Cheneval to delimit the scope of personal data to only include what the EU defines as personal data.
5. See e.g., (Narveson, 1999) for a good discussion of this proviso.
6. Given that Cheneval's account is supposed to be Lockean, it is not clear why only *partial* property rights obtain when someone mixes her labor with data about you. On Locke's account, mixing labor with something unowned does not only generate partial, but full property rights over the thing in question (provided a certain proviso is satisfied).
7. This is not a very Lockean idea. On Locke's account the first 'mixer' of labor would get full ownership, provided the proviso is satisfied. If Peter is the first to mix labor, then Carl and Allan get no ownership at all. But on Cheneval's account, where mixing of labor is a sufficient condition for at least partial ownership, Carl and Allan would indeed get at least partial ownership.
8. This example is inspired by a classic thought experiment from Kasper Lippert-Rasmussen's discussion of the SOT (Lippert-Rasmussen, 2008, p. 98).
9. Cheneval can still say that people can get partial ownership over personal data by engaging in certain contracts. For example, you could voluntarily engage in a contract with a company which allows them to mix their labor with the data and get partial ownership over it.

Acknowledgments

I would like to thank Beate Rössler, Rasmus Uhrenfeldt, Jørn Sønderholm and Jens Damgaard Thaysen for very useful comments on an earlier version of this paper.

Disclosure statement

No potential conflict of interest was reported by the author.

ORCID

Jakob Thrane Mainz  <http://orcid.org/0000-0002-7766-6439>

References

- Cheneval, F. (2018). Property rights of personal data and the financing of pensions. *Critical Review of International Social and Political Philosophy*, 1–23. <https://doi.org/10.1080/13698230.2018.1528521>
- Cohen, J. (2017). Law for the platform economy. *University of California, Davis Law Review*, 51(133), 133–204.
- Cwik, B. (2016). Property rights in non-rival goods. *Journal of Political Philosophy*, 24(4), 470–486. <https://doi.org/10.1111/jopp.12090>
- Floridi, L. (2013). *The ethics of information*. Oxford University Press.
- Laudon, K. (1996). Markets and privacy. *Communications of the ACM*, 39(9), 92–104. <https://doi.org/10.1145/234215.234476>
- Lippert-Rasmussen, K. (2008). Against self-ownership: There are no fact-insensitive ownership rights over one's body. *Philosophy & Public Affairs*, 36(1), 86–118. <https://doi.org/10.1111/j.1088-4963.2008.00125.x>
- Locke, J. (1690) [1988]. *Two treatises of government*. (P. Laslett, ed.). Cambridge University Press.
- Narveson, J. (1999). Property rights: Original acquisition and lockean provisos. *Public Affairs Quarterly*, 13(3), 205–227.
- Samuelson, P. (2000). Privacy as intellectual property? *Stanford Law Review*, 52(5), 1125–1173. <https://doi.org/10.2307/1229511>
- Thouvenin, F., Weber, R. H., & Früh, A. (2017). Data ownership: Taking stock and mapping the issues. In M. Dehmer (Ed.), *Frontiers in data science*. Frank Emmert-Streib, CRC Press.
- Union & regulation,. 2016 /67, article 4,1<https://eur-lex.europa.eu/eli/reg/2016/679/oj>.

INFERENCES AND THE RIGHT TO PRIVACY

In this paper, I defend what I shall call the ‘Inference Principle’. This principle holds that if an agent obtains some information legitimately, then it is legitimate for the agent to make any inference based on the information.¹ This principle is interesting for at least three reasons. First, it constitutes a novel answer to the timely question of whether the widespread use of ‘data analytics’ to infer personal information about individuals is morally permissible.² Second, it contradicts what seems to be a common view of inferences’ ability to violate privacy rights. Third, it offers an account of the theoretically underdeveloped issue of what duties are engendered by the moral right to privacy with regards to inferred information.³

State-of-the-art data analytics makes it possible to accurately infer all sorts of personal information about individuals, based on big data sets containing more or less trivial information, such as what car people drive, who their friends are, what groceries they buy, etc. Statistical correlations in the datasets reveal ‘new’ information about individuals, such as their political views, credit worthiness, or health conditions.⁴ The inferences are often used to develop machine learning models that predict the behavior of individuals. Political campaigns try to predict whom individual electors will vote for, banks try to predict if individuals will pay back a loan,⁵ insurance companies try to predict what health problems individuals will suffer from,⁶ authorities try to predict the risk of recidivism for individual prisoners who apply for parole,⁷ and Facebook and other tech

¹ By ‘legitimate’ I mean normatively legitimate, not epistemically legitimate. The epistemic version of this principle is often called ‘epistemic closure’, and holds, roughly, that if it is epistemically legitimate for Q to believe α and β , and α and β entail γ , then it is epistemically legitimate for Q to believe γ . See Luper 2020 for a good overview of the literature on epistemic closure.

² In recent years, there has been a legal discussion on whether inferences of personal information should be covered by the legal right to privacy (Wachter 2019; European Court of Justice 2017; Wachter & Mittelstadt 2019). To the extent that law should reflect morality, the Inference Principle has direct implications for this legal discussion.

³ Throughout this paper, I shall assume for the sake of argument that privacy rights exist. I shall not commit to any particular view on what a moral right in general consists in, or what the relation in general is between rights and duties. Neither shall I commit to any particular view on whether the right to privacy is an absolute right or not.

⁴ Barocas & Nissenbaum 2014, p. 44.

⁵ Turkson et al. 2016; Kearns & Roth 2020.

⁶ Price & Cohen 2019.

⁷ Berk & Hyatt 2015. See also Lin et al. 2020 for recent skepticism about the accuracy of these algorithms.

companies infer the preferences of individuals, in order to target them effectively with advertisement.⁸ Data analytics is used in a large variety of domains, and it influences more and more parts of our daily lives.

A common view in both philosophy, law, and computer science, is that the inferences of personal information infringe upon or violate individuals' privacy rights⁹, when the relevant individuals did not intend to disclose the inferred information.¹⁰ The idea seems to be that by training the machine learning models, personal information about individuals is accessed illegitimately, and that this is so, even if the inferences are based solely on publicly available information, or on information that the individual has shared voluntarily. In this journal, Benedict Rumbold and James Wilson have recently put this view as follows:

[...] we think that it is important to have an account of the right to privacy that at least makes it intelligible that such uses of information could violate privacy— that there can be cases in which an individual's right to privacy could be violated by the appropriation and dissemination of information either that they themselves have made public or that has been inferred from information they have made public.¹¹

Pace Rumbold and Wilson, the Inference Principle implies that an inference does not constitute a privacy violation, if the individual whom the information is about, has waived her right to privacy over the original information, on which the inference is based. Importantly, the Inference Principle does not imply that inferences cannot constitute privacy violations *simpliciter*. But since information is 'closed under entailment' – as

⁸ Tadesse et al. 2018.

⁹ There is no consensus in the literature on what the right to informational privacy is, and what counts as a violation of this right. So-called control theorists believe that an agent's right to privacy is violated when she loses the right kind of control over her personal information (or over the access to this information). For different versions of the control theory, see e.g. Moore 2003; Moore 2010; Inness 1992; Fried 1968; Parent 1983; Marmor 2015; Mainz & Uhrenfeldt 2020; Menges 2020. So-called access theorists often add the extra necessary condition that someone must actually access the agent's personal matters in order for her right to privacy to be violated. See e.g. Thomson 1975; Macnish 2018; Lundgren 2020. For present purposes, I shall remain agnostic about which of these theories, if any, is true. However, the argument I make in this paper may have revisionary implications for some of these theories.

¹⁰ Wachter 2019; Wachter & Mittelstadt 2019; Rumbold & Wilson 2019; Alben 2020; Barocas & Nissenbaum 2014; Kröger 2019.

¹¹ Rumbold & Wilson 2019, p. 3.

logicians say¹² - an individual who waives her right to privacy over some information also waives her right to privacy over any information that is inferred from it.¹³ Thus, if the Inference Principle is true, then it has implications for the moral permissibility of using data analytics to infer personal information about individuals. To wit, if the Inference Principle is true, then inferences of personal information constitute privacy violations far less often than we might think.

The paper proceeds as follows: In section I, I present and defend the Inference Principle. In section II and III, I present and reject two objections to my argument. Finally, in section IV, I make a few concluding remarks.

I. THE INFERENCE PRINCIPLE

According to the

Inference Principle: If an agent obtains some information legitimately, then it is legitimate for the agent to make any inference based on the information.¹⁴

If an agent Q obtains information α and information β legitimately, and γ can be inferred from α and β , then it is legitimate for Q to infer γ .¹⁵ Let us consider an example. Suppose that Smith tells Tom over the phone that Smith is on dialysis in his living room. Smith also sends Tom pictures of himself being connected to the dialysis machine. Tom is a medical doctor, and he knows that the only reason why one is on dialysis is that one has dysfunctional kidneys.¹⁶ Smith is unaware of this fact. Tom now makes the inference that

¹² Floridi 2006, p. 116.

¹³ To be clear, this is not to suggest that inferences cannot *diminish* an agent's privacy in a non-normative sense.

¹⁴ Note that the Inference Principle does not only involve 'personal' information. One reason for this is that it is notoriously difficult to distinguish personal information from non-personal information. A second reason is that pieces of information that are clearly personal can often be inferred from pieces of information that are clearly non-personal (Barocas & Nissenbaum 2014, p. 55). A third reason is that the principle also covers information that is completely non-personal in nature, regardless of where we draw the line between personal- and non-personal information.

¹⁵ The principle concerns *agents* in general, not only individuals. Nevertheless, throughout the paper, I will mostly talk about information about individuals, and inferences made by individuals.

¹⁶ For the sake of argument, set aside the off chance that Smith is on dialysis only because he likes it, has been forced to do it, or something similar.

Smith has dysfunctional kidneys. To make the inferences, Tom applies the standard logical inference rule of *conditional elimination* to α and β , and infers γ :

- (α) Smith is on dialysis.
- (β) If one is on dialysis, then one has dysfunctional kidneys.
- (γ) Smith has dysfunctional kidneys.¹⁷

Tom obtains the information that Smith is on dialysis in his living room legitimately. Plausibly, Tom obtains this information legitimately because Smith has waived his right to privacy over the information by intentionally disclosing the information to Tom.¹⁸ Tom also obtained legitimately the information that if one is on dialysis, then one has dysfunctional kidneys. The reason why it was legitimate for Tom to obtain this information is that Tom has read it in a standard medical textbook. According to the Inference Principle, it is legitimate for Tom to infer the information that Smith has dysfunctional kidneys. He simply applies a standard logical inferences rule to α and β in his mind, and infers γ . Thus, Tom does not violate Smith's right to privacy. This is so, even if Smith is unaware of the fact that the only reason why one is on dialysis is that one has dysfunctional kidneys, and even if Smith does not want Tom to know that he has dysfunctional kidneys. This seems to be an intuitively plausible result.

One can straightforwardly substitute $\{\alpha, \beta\}$ with any other set of legitimately obtained propositions containing information that is covered by the right to privacy, say, $\{\alpha_1, \beta_1\}$. Any correct inference to proposition $\{\gamma_1\}$ from the substituted propositions will

¹⁷ Conditional elimination is the inference rule at work in standard modus ponens arguments of the form 'if p then q, p, therefore q'. It makes no relevant difference what exact inference rule is at play. The reader can easily construct different inferences involving different inference rules.

¹⁸ There are two competing views in the literature on what it takes to waive one's right to privacy. The first view holds that the right to privacy is limited to information that the right-holder has not intentionally made public. For discussion of this view, see Thomson 1975; Reiman 1976; Fried 1968; Schoeman 1984; Parent 1983; Ryberg 2007. The second view holds that the right to privacy at least sometimes extends to information that the right-holder has intentionally made public. For discussion of this view, see Nissenbaum 1998, 2009; Stahl 2020; Timan et al. 2017; Roessler 2016; Newell et al. 2018; Moreham 2006; Reidenberg 2014; Rumbold & Wilson 2019; and Margulis 2003. For the purpose of this paper, I need not commit to a particular view on what is required to waive one's right to privacy. *Regardless* of what the correct view is, the Inference Principle implies that *if* an individual holds some information in accordance with this view, *then* the individual may legitimately infer any information from it. I remain non-committal about what is required in order to come to hold the original information legitimately.

also be legitimate.¹⁹ The principle also generates plausible results when the inference in question is not made in someone's mind but by, say, training a machine-learning model.

Consider an example. Jones owns a pickup truck, and Tim is the neighbor of Jones. Jones is proud of his car, and he frequently bores Tim with technical details about the car. Tim works as a data scientist. He wants to know what the correlations are between seemingly trivial data about electors, and their political preferences. He decides to find out whom Jones is likely to vote for in the upcoming election. He gets access to large amounts of data from publicly available databases, and trains a precise machine-learning model on the data. To his surprise, Tim discovers that owning certain types of pickup trucks is a very strong predictor of voting Republican, and that owning certain types of sedans is a very strong predictor of voting Democrat. Based on all the technical details about the car that Tim has listened to in the driveway, he knows that Jones owns the exact type of pickup truck that correlates very strongly with voting Republican. It so happens that Jones in fact always votes Republican. Jones does not want Tim to know his political preferences, and he is not aware that it is possible to infer his political preferences based on information about which car he drives.²⁰ Tim now asks the computer to calculate the likelihood of Jones voting Republican. Based on the correlations in the dataset, and the fact that Jones owns a specific type of pickup truck, the computer runs something like the following inference:

(α 1) Jones owns a pickup truck of type X.

(β 1) If one owns a pickup truck of type X, then one is very likely to vote Republican.

(γ 1) Jones is very likely to vote Republican.²¹

Tim obtains the information that Jones owns a pickup truck of type X legitimately. Plausibly, Tim obtains this information legitimately because Jones has waived his right to privacy over the information, by intentionally disclosing it to Tim. Tim also obtains

¹⁹ This is not to suggest that the Inference Principle only applies if one piece of information is inferred from two pieces of information. The number of members in the respective sets are not important. If, for instance, I hold the information legitimately that *all* men have a significant risk of getting testicular cancer, then I also hold legitimately that Smith has a significant risk of testicular getting cancer, Tom has a significant risk of getting testicular cancer etc.

²⁰ Car choice is in fact a good predictor of political preferences. Owners of pickup trucks are generally likely to vote Republican, and owners of sedans are generally likely to vote Democrat. See Gebru et al. 2017.

²¹ The model might output a precise estimation of the likelihood of Jones voting Republican. It might, for instance, output that Jones is 85% likely to vote Republican.

legitimately the information that owning a pickup truck of type X strongly correlates with voting Republican. The reason why it was legitimate for Tim to obtain this information is that he has obtained legitimately all the data necessary for training the machine-learning model. According to the Inference Principle, it is legitimate for Tim to use his computer to infer that Jones is very likely to vote Republican. Based on information that is legitimately obtained, he simply uses his computer to infer a ‘new’ piece of information about Jones. Thus, Tim does not violate Jones’ right to privacy. This is so, even if Jones was unaware of the fact that owning a pickup truck of type X correlates strongly with voting Republican, and even if Jones did not want Tim to know that he is very likely to vote Republican. Again, the Inference Principle generates an intuitively plausible result.

In the examples above, it is presumably epistemically permissible for Tom and Tim to make their respective inferences. By stipulation, the information on which the inferences are based, are obtained in epistemically legitimate ways. Because the original information has been broad about in epistemically legitimate ways, it is then epistemically legitimate for Tom and Tim to apply any valid logical inference rule to the information, and form a belief in the inferred information. To wit, all the steps that lead to Tom and Tim holding their respective inferred information are epistemically legitimate. If the Inference Principle was false, then it would be *epistemically legitimate* for Tom and Tim to make their respective inferences, but *morally illegitimate* to make said inferences. It remains contested whether epistemic duties and epistemic legitimacy are distinct from moral duties and moral legitimacy at all.²² But, even if these concepts are indeed distinct concepts, it seems strange that they should come apart in cases like the ones above, where *all* the steps that lead to the respective inferences are *both* epistemically legitimate, *and* morally legitimate. It would be strange that the very act of making the inference could be legitimate in one sense, yet illegitimate in the other, if there is nothing in the steps leading to the inference that can explain the difference in legitimacy.

We need not even look to epistemology to find the underlying idea that whatever arises from an unobjectionable situation by unobjectionable steps is itself unobjectionable. I am, of course, thinking of Robert Nozick’s famous entitlement theory of distributive justice here. According to Nozick, a distribution of goods cannot be unjust, if it arose from a just situation through a series of steps all of which were just.²³ Nozick criticized ‘end-state principles’ of justice, such as John Rawls’ Difference Principle, for being ahistorical.

²² See e.g. Wrenn 2007.

²³ Nozick 1974, p. 151.

He thought that in order to know whether a given distribution is just or unjust, we need to ask how the distribution came about.²⁴ If the distribution came about by through a series of just steps, from a distribution that is just, then the resulting distribution is just as well.

The Inference Principle resembles Nozick's point in the following way: According to the Inference Principle, it is legitimate to make an inference if the information that the inference is based on are obtained legitimately. We cannot simply ask the individual whom the information is about whether she wants the inferred information in question to be known by others. We cannot know whether the inference is legitimate without knowing how the inferred piece of information came about. If the inferred piece of information came about by making a correct inference²⁵ based on pieces of information all of which are obtained legitimately, then the inference is legitimate as well, even if the individual does not want the inferred information to be known by others. The basic idea of the Inference Principle is that if all steps in the process that leads to agent Q inferring information γ are legitimate, then it is difficult to see how it can suddenly be illegitimate for Q to infer γ . In the case of Smith and Tom, Tom obtains all the information relevant for making the inferences about Smith legitimately, and he makes the inference correctly. It is difficult to see how it then becomes illegitimate for Tom to infer the information that Smith has dysfunctional kidneys, given that all the steps that lead to Tom inferring this information were themselves legitimate.

Of course, Nozick's entitlement theory is controversial, and the Inference Principle might therefore be controversial as well. For any shortcoming the entitlement theory might have, we might worry that the Inference Principle inherits the same shortcoming. I will offer a few comments in mitigation of this worry.

Some of the well-known replies to objections against the entitlement theory also work for objections against the Inference Principle. Think for example of what we might call the 'Rectification Objection'. According to this objection, the distributions of many goods in the real world have historically *not* been distributed in accordance with the entitlement theory, and that current distributions of these goods are therefore unjust.²⁶ The

²⁴ Nozick 1974, p. 153-155.

²⁵ If the inference was not made correctly, then it might have generated a false belief in Tom's mind. Theorists who follow Prosser's theory of the right to privacy might argue that this would violate Smith's right to privacy (Prosser 1960, p. 389). I find it strange, though, that producing false beliefs about other people should violate their right to privacy, but for the sake of argument, I simply stipulate that Tom makes the correct inference.

²⁶ See Nozick's own discussion of this objection in Nozick 1974, p. 152-153.

corresponding objection against the Inference Principle holds that the ways in which private companies and governments in the real world have acquired individuals' personal information are illegitimate, and therefore the inferences they draw from them are illegitimate as well. Nozick's reply to the Rectification Objection against the entitlement theory is to concede that many goods should indeed, one way or another, be redistributed to their legitimate owners, or that at least the individuals who are worse off due to the historical injustices should somehow be compensated.²⁷ A similar reply works to the corresponding objection to the Inference Principle. Although endorsing the Inference Principle does not imply this, I can simply concede that much of the personal information that real life inferences are based on, are obtained illegitimately, and that the inferences based on them are thus illegitimate as well.

Even though the Inference Principle does not imply that inferences based on illegitimately obtained information are themselves illegitimate, the principle *does*, however, imply the following by contraposition: *If γ is obtained illegitimately, then either α or β - from which γ is inferred - is also obtained illegitimately.* If making the inference constitutes a further violation of the right to privacy, then it is because a violation already occurred in the process leading up to the inference. But, the Inference Principle neither implies that the inference *does* constitute a violation of the right to privacy, nor that it *does not* constitute a violation of the right to privacy. It is simply silent on the matter.

Here is yet another reason why the Inference Principle is not vulnerable to the classic objections against the entitlement theory: Even very unequal distributions of personal information do not generate the same intuitions of injustice as very unequal distributions of primary goods do to many people. If agent Q comes to hold a lot of personal information about agent P, while agent R holds no personal information about P, then - under normal circumstances - this does not generate the intuition that the distribution of personal information is unjust. It is perfectly consistent to be an egalitarian with respect to primary goods, while still endorsing the Inference Principle. Even the Lockean proviso does not apply when it comes to personal information. It would be strange to claim that the fact that Q acquires a certain amount of personal information about P is unjust because it does not leave 'enough and as good' for R. Even if the proviso did apply, it would be easily

²⁷ See Nozick 1974, p. 228-231.

satisfied given that information is generally a non-rivalrous good, which Q can have and use without preventing R from doing the same, and vice versa.²⁸

Before closing this section, let me offer two additional reasons for why the Inference Principle is plausible. The first reason is that if Tom violates Smith's right to privacy by making the inference about Smith's medical condition, then it implies that having certain thoughts in one's mind can – in itself - constitute rights violations. But this is very controversial. Many theorists maintain that having thoughts in one's mind does simply not seem to be the type of action that can constitute *rights violations*. One may have certain racist thoughts, but merely having these thoughts does not violate the rights of anyone. If these racist thoughts cause one to perform conduct that discriminate against members of a certain race simply because they are members of that race, then the discriminating conduct may constitute rights violations. But, having the thoughts that caused one to perform the conduct does not *in itself* constitute a rights violation.²⁹ I think this view is correct, but my argument does not rest on it. If the reader believes that having certain thoughts in one's mind can be wrongful or even violate the rights of others, then it does not undermine my argument. To see this, note that in the example of Jones and Tim, no inference is made in the mind of Tim. He merely lets the machine-learning model do all the work for him, and then only looks at the output data of the algorithm. Still, he does not violate Jones' right to privacy. If Tim comes to believe that Jones will vote Republican, but Tim has no idea how the algorithm reached this result, it would be strange to hold that Tim has now violated Jones' right to privacy.

This idea is reflected in the often-cited assumption in the privacy literature that 'simply knowing' something about an individual is not sufficient to violate the individual's right to privacy. Judith Jarvis Thomson, for instance, writes:

I should say straightaway that it seems to me none of us has a right over any fact to the effect that that fact shall not be known by others. You may violate a man's

²⁸ Mainz 2020, p. 5.

²⁹ Some authors do indeed seem to think that having certain thoughts can be harmful to others, because of downstream consequences caused by the thoughts (See Mendlow 2018; Dan-Cohen 1999; Morris 1976). Others believe, perhaps controversially, that having certain thoughts can be wrongful in itself, despite the lack of any upstream or downstream explanations (Schroeder & Basu 2018). Schroeder & Basu touch upon the idea that having certain *beliefs* about others may violate their right to privacy (Schroeder & Basu 2018), but the standard view seems to be that beliefs cannot constitute rights violations.

right to privacy by looking at him or listening to him; there is no such thing as violating a man's right to privacy by simply knowing something about him.³⁰

This assumption has been echoed by many others.³¹ The assumption essentially holds that we do not have privacy duties to not have certain beliefs about others, or at least that simply having a belief (even if the belief amounts to knowledge) about someone cannot *in itself* amount to a failure to comply with a duty strong enough that it constitutes a violation of a moral right to privacy.³² Following the literature, I shall assume that this is a plausible assumption.

The assumption is relevant for the plausibility of the Inference Principle for the following reason: If simply knowing something about an individual cannot in itself constitute a violation of her right to privacy, then it is all the more plausible that for any conduct that constitutes a violation of the right to privacy, the conduct must occur in the process that leads to the formation of knowledge about the individual. And, if this process consists only of steps all of which are legitimate, then it is all the more difficult to see where the wrongness that makes up the violation comes from. If the mere fact that Tom knows γ about Smith cannot constitute a violation of Smith's right to privacy, and the way in which Tom obtained α and β , from which γ is inferred, is legitimate, then it is difficult to see how Tom violates Smith's right to privacy.³³

³⁰ Thomson 1975, p. 307.

³¹ See e.g. Marmor 2015; Kappel 2013; Persson & Savulescu 2019. See, however, Munch 2021a for a critical discussion of this assumption.

³² Plausibly, we do, however, sometimes have duties to have certain beliefs about others. However, many of these duties are explained by their downstream consequences. One may for instance have a doxastic duty to have a certain belief, if forming this belief is necessary to perform an action that one has a duty to perform. To illustrate, a medical doctor who has a duty to treat a patient has an appertaining doxastic duty to form a belief about, say, what disease the patient suffers from. Similarly, one may have a doxastic duty *not* to form certain beliefs, if not forming such beliefs is necessary to perform an action that one has a duty to perform. The medical doctor may have a doxastic duty not to form the belief that the patient suffers from a disease that she does not suffer from. However, failure to comply with doxastic duties like these does not, in itself, constitute a violation of a moral right to privacy.

³³ Note that the on some views on the justification of privacy rights, the explanation for why some steps that lead to Q holding α , β , or γ are illegitimate have to do with the *consequences* of Q holding α , β , or γ . For instance, some privacy scholars think that the right to privacy is explained by an urgent moral interest in exercising control over how we present ourselves to others (See Marmor 2015). Other privacy scholars think the right to privacy is explained by an interest in avoiding that our personal information is somehow misused or exploited (See Parent 1983; Munch 2020), or because others' access to our personal information somehow detracts our ability to autonomously form

This concludes my positive defense of the Inference Principle. Even if my positive argument for the Inference Principle is unsatisfying, I believe that the arguments for the negation of the Inference Principle are even less satisfying. In the following sections, I present two objections to my argument that each gives reason to think that the negation of the Inference Principle is true. If I succeed in refuting these objections, then we have good reason to think that the Inference Principle is true.³⁴

II. THE INTENTIONALITY OBJECTION

Rumbold and Wilson argue that just because P has intentionally made α and β public, and Q infers γ from α and β , it does not mean that P has waived her right to privacy over γ with regards to Q. Whether P has waived her right to privacy over γ depends on whether P intended γ to be public as well, when P intentionally made α and β public.³⁵ Simply put, Rumbold and Wilson believe that the waiving of privacy rights over information tracks intentionality.³⁶ The right to privacy over a piece of information is waived if, and only if, the claimant intended that piece of information to be public, regardless of whether the information is inferred from some other information.

Rumbold and Wilson begin their argument with a critique of Thomson. Thomson claims that the right to privacy does not cover information that one has voluntarily disclosed. To illustrate her view, she gives the following example: Suppose you own a picture of yourself. You have a right that others do not look at the picture. Now consider the following options you have with regards to your picture and other people's access to it. You might

our identities, or detracts our ability to make autonomous decisions (See Feinberg 1986; Taylor 2002).

³⁴ Of course, there may be other objections to my argument. One candidate might be derived from the view recently defended by Lauritz Munch (2021b). He defends what he calls the 'symmetry thesis'. According to this thesis, there are no good reasons to think that there are any privacy-related normative differences between standard cases where someone accesses someone else's information by using an X-ray device, and cases where the exact same information is accessed through the means of statistical inferences. It is beyond the scope of this paper to provide a satisfying reply to Munch's argument. However, I think that the Inference Principle offers a plausible explanation for why we often find X-ray cases objectionable, and statistical cases unobjectionable: If the information that the inference is based on are obtained legitimately, then the inference does not constitute a privacy violation.

³⁵ For a similar point, see Floridi 2006, p. 116.

³⁶ Rumbold & Wilson 2019, p. 12.

- (1) invite others to look at it,
- (2) get others to look at it whether they want to or not,
- (3) let others look at it,
- (4) absentmindedly leave it somewhere where others would have to go through some trouble to look at it, or
- (5) absentmindedly leave it somewhere where nobody could reasonably be expected to know that it was owned by someone.³⁷

According to Thomson, you have waived your right to privacy in (1)-(5). In (1), (2), and (3), the right is waived *intentionally*, and in (4) and (5) it is waived *unintentionally*. Thomson's view captures the intuition that if I, for instance, walk down the street, then other people do not violate my right to privacy when they look at me. They may get access to all sorts of information about me, like information about what clothes I wear, what physical disabilities I have, etc. But, because I intentionally make this information public by walking down the street knowing that others can easily get access to the information, it seems strange to claim that they now violate my right to privacy when they look at me.

Rumbold and Wilson argue, *pace* Thomson, that it is generally impossible to waive one's rights unintentionally:

In particular, it seems odd to claim that one could waive a right unintentionally. Rather, if one is to waive a right, one would seem to need actually to waive it—the very notion of ‘waiving’ implying an intentional action on the part of the relevant agent with regard to their right.³⁸

Thus, Rumbold and Wilson believe that the right to privacy actually covers (4) and (5). Supposedly, it is the absentmindedness of the right-holder that leads Thomson to conclude that the right to privacy is waived in (4) and (5). But, as Rumbold and Wilson point out, absentmindedness normally entails neither waiving nor forfeiture of rights. Just because you absentmindedly leave the car keys in your car, it does not mean that you have waived

³⁷ Thomson 1975, p. 301.

³⁸ Rumbold & Wilson 2019, p. 10.

or forfeited your property rights over the car.³⁹ If someone drives away in the car, he is a car thief and not the happy owner of a new car.

The next step in Rumbold and Wilson's argument is to claim that Thomson's logic must also "... cover anything anyone might infer from looking at the picture".⁴⁰ For example, they say, if the picture is of you in high school, someone might be able to infer, with a varying degree of accuracy, which school you went to, how happy you were at that time etc.⁴¹ If you have a right that others do not look at the picture, then presumably you also have a right that they do not infer any information from looking at the picture. If so, then presumably you also waive your right to privacy over the inferred information, when you waive it over the picture. Thus, according to Rumbold and Wilson, Thomson's view entails the Inference Principle – although they do not use this terminology.

Rumbold and Wilson find the Inference Principle implausible. They support their view by use of the following hypothetical:

Imagine Annabel. Annabel is a famous actress. She also suffers from a rare and very hard to diagnose genetic disorder, a piece of information about herself she wishes to keep private. One day, Annabel agrees to take part in a new medical initiative. The primary purpose of the initiative is to promote the donation of genetic code for research purposes. As a participant in the initiative, Annabel agrees to donate her DNA to medical science and, to allay the public's worries about genetic research, even agrees to post it on the internet, together with a note advertising the fact that it is hers. Unbeknownst to Annabel, however, by posting this information on the internet, Annabel also makes it possible for those trained in genetic medicine to deduce that she suffers from her rare genetic disorder. Brian is one such researcher and, having studied Annabel's DNA, decides to go to the papers to publicize that fact.⁴²

If Rumbold and Wilson's view is correct, then Brian violates Annabel's right to privacy. The information about her genetic disorder was inferred from the public DNA profile, but Annabel did not intend to make the information about the genetic disorder public. So, Brian

³⁹ Rumbold & Wilson 2019, p. 15.

⁴⁰ Rumbold & Wilson 2019, p. 4.

⁴¹ Rumbold & Wilson 2019, p. 14.

⁴² Rumbold & Wilson 2019, p. 14.

violates Annabel's right to privacy by making the inference, according to Rumbold and Wilson. Rumbold and Wilson in effect treat the Annabel case as a counterexample to the Inference Principle. So, if Rumbold and Wilson's view is true, then the Inference Principle is false.

I think that there are good reasons to reject Rumbold and Wilson's view. The first reason is that it rests on a questionable assumption. Recall that Rumbold and Wilson assume that it is generally impossible to waive a right *unintentionally*. But, this assumption is not nearly as obvious as Rumbold and Wilson seem to think. As Lauritz Munch have recently pointed out, some accounts of consent imply that it is indeed possible to waive a right unintentionally:

Their argument [Rumbold and Wilson's, red.] relies on appealing to the thought that rejecting their view allows for cases in which people would have waived their (privacy) rights without doing so intentionally, which they deem theoretically problematic. However, it is not clear what precisely is the theoretical cost of accepting the possibility of some such cases. Plausibly, any account of consent under which consent is an act of communication must allow that there is sometimes a disconnect between people's intentions and the communicative act that validly instantiates the consent [...].⁴³

Communication accounts of consent allow for unintentional waivers of rights, at least in some cases. On such accounts, the right-holder's intentions can be misaligned with what is actually communicated.⁴⁴ If Smith by his own actions communicate consent to Tom accessing his medical information, then Smith has plausibly waived his right to privacy over this information with regards to Tom, even if Smith never intended to do so. On such accounts of consent, it is perfectly consistent to hold – as the Inference Principle implies – that it possible to unintentionally waive one's right to privacy over some inferred piece of information, if one has waived one's right to privacy over the information on which the inference is based. It thus seems theoretically uncstly to reject the assumption on which Rumbold and Wilson rest their argument.

⁴³ Munch 2021b, p. 4.

⁴⁴ However, as Munch notes, it is presumably desirable to minimize the occurrences of such misalignments. See Bolinger 2019 for discussion of this.

Now, Rumbold and Wilson are aware that their view has the strange implication that making an inference from legitimately obtained information can violate the right to privacy. They write:

However, it is also clear that at this point our model faces certain difficulties. For example, imagine that, rather than posting her DNA on the internet, during a party Annabel happens to bump into Sherlock Holmes. As the world knows, Holmes is a master of both observation and deduction and, during their conversation, he is able to deduce by mentally interrogating a series of stories Annabel tells him that she suffers from the rare genetic condition that she has tried so desperately to keep private. What kind of duties might Holmes be under at this point? On our model, it is not just that Holmes is under a duty to refrain from publicizing Annabel's condition but, perhaps more surprisingly, that he infringes (possibly even violates) Annabel's right to privacy insofar as he makes any effort to deduce the nature of Annabel's condition in the first place (to 'appropriate' it from information Annabel makes public).⁴⁵

This is indeed a difficulty for Rumbold and Wilson's view, and I think they underestimate the degree to which this is so. It seems odd that not only does Sherlock violate Annabel's right to privacy if he publicizes the information about Annabel's disorder, he also violates (or at least infringes upon) her right to privacy by simply making the inference. Rumbold and Wilson's solution to this problem is to 'bite the bullet':

In those cases, then, where we, as duty-bearers, know that a piece of once private information is private and that the relevant right-bearer has only made it public unintentionally, we find ourselves ready to bite the bullet. That is, insofar as P has been attempting to keep a given piece of information private and we, as duty-bearers, know this, we believe that it would infringe her right to privacy were we to appropriate it by inferring it from information she has made public (intentionally or not).⁴⁶

⁴⁵ Rumbold & Wilson 2019, 13.

⁴⁶ Rumbold & Wilson 2019, p. 14.

Given that Rumbold and Wilson acknowledge that their view has a strange implication, it is puzzling that they do not seem to consider the further implications of simply biting the bullet. In relation to the case involving Tim and his inference of Jones' political preference, Rumbold and Wilson's view implies that Tim remains under an obligation not to make the inference about Jones' political preference, if Tim knows that Jones did not wish his political preference to be public. We have already seen that it is theoretically uncult to drop the assumption that it is impossible to waive a right unintentionally. But, even if we keep this questionable assumption for the sake of argument, Rumbold and Wilson's view still has a strange implication: Suppose that Tim knows the correlations between owning a certain type of pickup truck, and having certain political preferences, long before Jones moves in. Suppose further that Tim knows that he would not be able to stop himself from making the inference about Jones' political preference, had he known which car Jones drives.⁴⁷ If Rumbold and Wilson's view is correct, then presumably Tim now has at least a *pro tanto* obligation to make an effort to avoid knowing which car Jones drives. When Jones comes home, Tim has an obligation to look away before he sees Jones' car. When Jones starts talking to Tim in the driveway about his new car, Tim has an obligation to put his fingers in his ears or otherwise prevent Jones from telling him what car it is.⁴⁸ Note that it does not even matter if Jones really *wants* Tim to know what car he drives. He can intentionally waive his right to privacy over information about what car he drives, and Tim still violates Jones' right to privacy the second he receives the information and makes the inference, according to Rumbold and Wilson's view. Similarly, if Tim has access to all the relevant information, but simply lacks the logical reasoning skills necessary to make the inference, then Tim now has at least a *pro tanto* obligation not to take an introductory logic course, or otherwise engage in activities that could teach him how to make the inference, if he knows that he would not be able to refrain from making the inference if he knew how

⁴⁷ I presume that as a matter of psychological fact, it is at least sometimes *impossible* to form a certain belief *b* (or refrain from forming *b*) at will. This view, or at least something close to it, is known as 'Doxastic Involuntarism'. See Peels 2015; Antill 2020; Roeber 2019. I also presume that it is at least sometimes psychologically *possible* to know in advance that one would not be able to refrain from forming *b* if one was presented with evidence *e*.

⁴⁸ Munch has recently called a duty of this type an 'indirect doxastic duty' not to form a certain belief (Munch 2021a). It is an *indirect* doxastic duty because the duty consists in acting in a way that indirectly avoids forming the belief in question.

to make it.⁴⁹ He would have an obligation to remain ignorant of how the inference rule of *conditional elimination* works.

Now, Rumbold and Wilson might respond by saying that Tim indeed has a *pro tanto* obligation to remain ignorant about what car Jones drives, but that this obligation is rendered defunct because it would then be too demanding for Tim to comply with Jones' right to privacy. This response is available to Rumbold and Wilson, but it is a response that sits uncomfortably with their view of how strong the right to privacy is. Rumbold and Wilson seem to believe that privacy interests are so important that your privacy interests even have to be protected against other people inferring - in their minds - relatively trivial information about you. But if these interests are so important, then it seems strange that Tim's obligation not to violate Jones's right to privacy is rendered defunct when Tim is able to infer *non-trivial* information about Jones' political preferences. If Jones' privacy interests are so important, then one should expect that Tim has an obligation to make significant efforts to avoid violating Jones' privacy rights.

Rumbold and Wilson might then respond by saying that they explicitly acknowledge that if Jones gives Tim information $\alpha 1$ and $\beta 1$, and Tim cannot help but to infer $\gamma 1$, then Tim's duty not to infer $\gamma 1$ is rendered *defunct*, and therefore Tim does not violate Jones' right to privacy when he infers $\gamma 1$.⁵⁰ If ought implies can, and Tim cannot avoid making the inference, then Tim does not have a duty not to make the inference. However, the situation involving Tim and Jones is different. Tim knows *beforehand* that we will not be able to avoid making the inference, if he is presented with evidence of what car Jones drives. In this case, Tim actually *can* avoid making the inference, so his duty is *not* rendered defunct on Rumbold and Wilson's view. Thus, this response is not available to Rumbold and Wilson.

Yet another - and possibly worse - problem with Rumbold and Wilson's view is that it (in contrast to the Inference Principle, as we have seen) has the controversial implication that having certain thoughts in one's mind can in itself constitute rights

⁴⁹ Or, he might be under an obligation to *become* ignorant of the information about Smith that he already knows. Becoming ignorant of information that one already knows may be psychologically possible in epistemically non-drastic ways in at least some cases (Matheson 2013). But it is still normatively controversial to hold that one can be under an obligation to become ignorant of certain information. In the case of Jones and Tim, it would be very strange to claim that Tim has an obligation to become ignorant about either the information about what car Jones drives, or the information about statistical correlations between car choice and political preferences, given that Tim has come to know both pieces of information in legitimate ways. Jones even *wants* Tim to know what car he drives.

⁵⁰ Rumbold & Wilson 2019, p. 14.

violations. Rumbold and Wilson think that Sherlock violates Annabel's right to privacy by making inferences in his mind, and thus they believe that having certain thoughts in one's mind can violate the rights of others. This implication is very controversial, although as we saw in an earlier section, some theorists would be willing to bite the bullet here. However, the burden of proof still seems to be on Rumbold and Wilson, because their argument has an implication that has the very controversial implication that making inferences in one's mind can constitute a rights violation.

Now that we have seen why Rumbold and Wilson's Intentionality Objection is mistaken, let us now turn to what I call the 'Other-Regarding Inference Objection'. Albeit being very common, this objection nevertheless turns out to be mistaken too.

III. THE OTHER-REGARDING INFERENCE OBJECTION

The Other-Regarding Inference Objection holds that legitimately obtained information about some individual(s) can lead to illegitimate inferences about *others*. The objection comes in several versions, but I shall focus on the version that threatens the Inference Principle the most.⁵¹ The reader may be familiar with something like the following scenario: You open your laptop and go online. To your surprise, you see advertisement for your favored political party on all the websites you visit, despite the fact that you do your best to hide your political preferences online. You do some research, and discover that Facebook has – even though you do not have a Facebook profile yourself - inferred your political preference from information about your friends' political preferences that they have voluntarily shared online, in combination with the publicly available information that you are friends with them, and the publicly available information that friends often share political preferences.⁵² You feel that Facebook violates your right to privacy by inferring your political preference. Facebook may have obtained the information about your friends' political preferences legitimately, they may have obtained legitimately the information that you are friends with them, and, they may have obtained legitimately the information that friends often have the same political preferences. But – the objection goes – it is not legitimate to infer *your* political preference legitimately, because you did not contribute to the information from which the inference was made.

⁵¹ Another version of this objection can be found in Floridi 2006, p. 116.

⁵² Facebook have created so-called 'shadow profiles' of people who do not have a Facebook profile. These profiles also contain inferred information about non-users based on information about users, and certain connections between users and non-users (Garcia 2017).

This scenario is relevantly different from the scenario of Smith and Tom, and the scenario of Jones and Tim. In those scenarios, the right-holders *did* provide some of the information from which the inferences were made. In the case of Smith and Tom, Smith voluntarily shared with Tom the information that he was on dialysis in his living room. In the case of Jones and Tim, Jones voluntarily shared with Tim the information that he owned a pickup truck. This difference in the voluntary sharing of some of the original information makes a morally relevant difference. Or so the objection goes.

I do not think that Facebook violates your right to privacy when they infer your political preference. If Facebook's behavior counts as a violation of your right to privacy, then we all go around violating each others' privacy rights all the time in the analog world. Almost all of the information we voluntarily share with others all the time can be used to infer information about third parties.⁵³ It makes no principled difference that Facebook makes inferences about *many* individuals, or that the Facebook's inferences are *sometimes* more accurate. Suppose that individual P tries to hide her political preference A, while her friends are very outspoken about their own political preference A. Suppose further that groups of friends are in fact very likely to have the same political preferences.⁵⁴ I know P and her friends, and I now make the inference that P has preference A. We all make similar inferences of personal information about other people. But, we do normally not think that doing so amounts to violations of their privacy rights. If making such inferences violate the privacy rights of others, then this suggests an extremely revisionary theory of privacy rights.

I do not suggest that Facebook is not acting wrongly, all things considered. If Facebook is acting wrongly in the scenario above, then it might be because they *use* the inferences in illegitimate ways. Perhaps the 'micro-targeting' of political advertisement amounts to a problematic form of voter manipulation.⁵⁵ Or, perhaps the inferences lead Facebook's algorithms to distribute fake news and conspiracy theories to voters who are likely to believe them and vote accordingly. However, whatever may be wrong with the way the inferences are used, I do not think that merely making the inference – in itself – is illegitimate. In particular, I do not think that it violates anyone's right to privacy.

Even if I am mistaken, and Facebook's inference of your political preference *does* violate your right to privacy, the Inference Principle can in fact handle this. Recall

⁵³ Floridi 2006, p. 116.

⁵⁴ This is in fact true in many cases. See Pew Research Center 2014.

⁵⁵ See Susser et al. 2019 for a discussion of this view.

Nozick's reply to the Rectification Objection. The reply was simply to concede that many goods should indeed, one way or another, be redistributed to their legitimate owners, or, at least the individuals who are worse off due to the historical injustices should somehow be compensated. Again, I can concede something similar. I can concede that Facebook violates your right to privacy in the process that leads to the inference, and that the inference is therefore illegitimate. One of the pieces of information that makes it possible to infer your preference, is the information that you are friends with individuals who have that preference. The way in which Facebook have historically gathered information like this in the real world is questionable at best. Facebook asks users if they want to 'import their friends' from their phones, to make it easier to connect with their friends. The potential problem is that by consenting to this, users give Facebook access to their friends' names, contact information etc., without consent from the friends.⁵⁶ Based on the information, Facebook then infers personal information about the friends, even if they do not themselves have a Facebook profile.⁵⁷ Thus, even if Facebook violates your privacy by making the inference, I can simply concede that this is so exactly because some of the original information is obtained illegitimately, and not because the Other-Regarding Inference Objection is true.

IV. CONCLUDING REMARKS

In this paper, I have defended the Inference Principle. If this principle is correct, then it has wide-reaching implications for the moral permissibility of inferring personal information by using data analytics. So far, many commentators have claimed that the inferences of personal information violate people's privacy rights. The Inference Principle implies that this is not always so. At least in cases where the inferences are based solely on information that is obtained legitimately, the inferences do not violate privacy rights. This result is of theoretical philosophical interest, but it also suggests that the use of data analytics is morally permissible in a surprisingly wide range of cases. The Inference Principle even offers concrete action guidance to data analysts: If all the information in a database is obtained legitimately, then the data analyst is morally permitted to make any inference from the information. The Inference Principle also offers concrete advice to policy makers: If they want law to track morality, then privacy laws and data protection regulations should

⁵⁶ Garcia 2017.

⁵⁷ Garcia 2017, p. 1.

probably not be extended to cover inferences made from information that is obtained legitimately.

Bibliography

- Alben, Alexander. 2020. "When Artificial Intelligence and Big Data Collide—How Data Aggregation and Predictive Machines Threaten our Privacy and Autonomy". *AI Ethics Journal*. 1(1): 1-23.
- Antill, Gregory. 2020. "Epistemic Freedom Revisited," *Synthese*. 197: 793–815.
- Barocas, Solon. Helen Nissenbaum. 2014. "Big Data's end run around anonymity and consent". *Privacy, Big Data, and the Public Good: Frameworks for Engagement*, ed. Julia Lane, Victoria Stodden, Stefan Bender, and Helen Nissenbaum. *Cambridge: Cambridge University Press*.
- Berk, Richard. Jordan Hyatt. 2015. "Machine Learning Forecasts of Risk to Inform Sentencing Decisions". *Federal Sentencing Reporter*. 27(4): 222–28.
- Bolinger, Renée. 2019. "Moral risk and communicating consent." *Philosophy and Public Affairs*. 47: 179–207.
- Dan-Cohen, Meir. 1999. "Harmful Thoughts". *Law and Philosophy*. 18(4): 379-405.
- European Court of Justice. 2017. *Peter Nowak v Data Protection Commissioner* Case C-434/16.
- Feinberg, Joel. 1986. *The Moral Limits of the Criminal Law, vol. 3: Harm to Self*. Oxford: Oxford University Press.
- Floridi, Luciano. 2006. "Four Challenges for a Theory of Informational Privacy". *Ethics and Information Technology*. 8: 109-119.
- Fried, Charles. 1968. "Privacy". *Yale Law Journal*. 77(3).
- Garcia, David. 2017. "Leaking Privacy and Shadow Profiles in Online Social Networks." *Science Advances*. 3(8).
- Geburu, Timnit, Jonathan Krause, Yilun Wang, Duyun Chen, Jia Deng, Erez Lieberman Aiden, and Li Fei-Fei. 2017. "Using deep learning and Google Street View to estimate the demographic makeup of neighborhoods across the United States." *Proceedings of the National Academy of Sciences*. 114(50):13108.
- Kappel, Klemens. 2013. "Epistemological Dimensions of Informational Privacy." *Episteme* 10(2): 179-192.
- Kearns, Michael. Aaron Roth. 2020. *The Ethical Algorithm*. Oxford: Oxford University Press.
- Kröger J. 2019. "Unexpected Inferences from Sensor Data: A Hidden Privacy

- Threat in the Internet of Things.” *IFIP Advances in Information and Communication Technology*. Vol. 548.
- Lin, Zhiyuan. Jongbin Jung. Shared Goel. Jennifer Skeem. 2020. “The limits of human predictions of recidivism”. *Science Advances* 6(7).
- Luper, Steven. 2020. "Epistemic Closure". *The Stanford Encyclopedia of Philosophy*. URL =<<https://plato.stanford.edu/archives/sum2020/entries/closure-epistemic/>>.
- Macnish, Kevin. 2018. “Government Surveillance and Why Defining Privacy Matters in a Post-Snowden World.” *Journal of Applied Philosophy*. 35(2): 417–432.
- Mainz, Jakob. 2020. “But Anyone Can Mix Their Labor: A Reply to Cheneval”. *Critical Review of International Social and Political Philosophy*. Forthcoming.
- Mainz, Jakob. Rasmus Uhrenfeldt. 2020. “Too Much Info: Data Surveillance and Reasons to Favor the Control Account of the Right to Privacy.” *Res Publica*. Forthcoming.
- Margulis, S. T. 2003. “Privacy as a Social Issue and Behavioral Concept”. *Journal of Social Issues*. 59(2): 243-261.
- Marmor, Andrei. 2015. “What is The Right to Privacy?” *Philosophy & Public Affairs*. 43(1): 3-26.
- Matheson, David. 2013. “A Duty of Ignorance”. *Episteme*. 10(2): 193–205.
- Mendlow, Gabriel. 2018. “Why Is It Wrong To Punish Thought?” *The Yale Law Journal*. 127(8): 2204-2585.
- Menges, Leonhard. 2020. “A Defense of Privacy as Control.” *Journal of Ethics*. <https://doi.org/10.1007/s10892-020-09351-1>
- Moore, Adam D. 2003. Privacy: Its Meaning and Value. *American Philosophical Quarterly* 40(3): 215–227.
- Moore, Adam D. 2010. *Privacy Rights: Moral and Legal Foundations*. University Park: Pennsylvania State University Press.
- Moreham, N. 2006. “Privacy in Public Places”. *Cambridge Law Journal*. 65(3): 606-635.
- Munch, Lauritz. 2020. “The Right to Privacy, Control Over Self-Presentation, and Subsequent Harm.” *Journal of Applied Philosophy*. 37(1): 141–154.
- Munch, Lauritz. 2021a. “How Privacy Rights Engender Direct Doxastic Duties.” *The Journal of Value Inquiry*. Forthcoming.
- Munch, Lauritz. 2021b. “Privacy Rights and ‘Naked’ Statistical Evidence.” *Philosophical Studies*. Forthcoming.
- Newell, B. I. Skorvanak. T. Timan. T. Chokrevski. 2018. “A Typology of Privacy.”

- University of Pennsylvania Journal of International Law*. 38(2): Art. 4.
- Nissenbaum, Helen. 1998. "Protecting privacy in an information age: the problem of privacy in public". *Law and Philosophy*. (17): 559–96.
- Nissenbaum, Helen. 2009. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford, CA: Stanford University Press.
- Nozick, Robert. 1974. *Anarchy, State, and Utopia*. Basic Books.
- Parent, William. 1983. "Privacy, Morality, and Law". *Philosophy & Public Affairs*. 12(4): 269-288.
- Peels, Rik. 2015. "Believing at Will Is Possible." *Australasian Journal of Philosophy* 93(3): 524–541.
- Persson, Ingmar. Julian Savulescu. 2019. "The Irrelevance of a Moral Right to Privacy for Biomedical Moral Enhancement," *Neuroethics*. 12(1): 35–37.
- Pew Research Center. 2014. "Political Polarization in the American Public." URL=<https://www.pewresearch.org/politics/2014/06/12/political-polarization-in-the-american-public/> (Accessed 21-01-2021).
- Price, W.N. I.G. Cohen. 2019. Privacy in the age of medical big data. *Nature Medicine* 25, 37–43.
- Prosser, William. 1960. "Privacy". *California Law Review*. 48(3): 383-423.
- Reidenberg, Joel. 2014. "Privacy in Public". *University of Miami Law Review*. 69(1).
- Reiman, Jeffrey. 1976. "Privacy, Intimacy, and Personhood". *Philosophy & Public Affairs*. 6(1): 26-44.
- Roeber, Blake. 2019. "Evidence, Judgment, and Belief at Will." *Mind*. 128(511): 837–859.
- Rumbold, Benedict. James Wilson. 2019. "Privacy Rights and Public Information." *The Journal of Political Philosophy*. 27(1): 3–25.
- Roessler, Beate. 2016. "Privacy as a Human Right." *Proceedings of The Aristotelian Society*. CXVII(2).
- Ryberg, Jesper. 2007. "Privacy Rights, Crime Prevention, CCTV, and the Life of Mrs. Aremac." *Res Publica*. 13: 127-143.
- Schoeman, Ferdinand. 1984. *Philosophical Dimensions of Privacy: An Anthology*. Cambridge University Press.

- Schroeder, Mark. Rima Basu. 2018. "Doxastic Wronging". In *Pragmatic Encroachment in Epistemology*, Brian Kim, Matthew McGrath (eds.). Routledge.
- Stahl, Titus. 2020. "Privacy in Public: A Democratic Defense". *Moral Philosophy and Politics*. 7(1): 73-96.
- Susser, Daniel. Beate Roessler. Helen Nissenbaum. 2019. "Online Manipulation: Hidden Influences in a Digital World." *Georgetown Law Technology Review*. 4(1).
- Tadesse, Michael M. Hongfei Lin. Bo Xu. Liang Yang. 2018. "Personality Predictions Based on User Behavior on the Facebook Social Media Platform". *IEEE Access* vol. 6, 61959-61969.
- Taylor, James Stacey. 2002. "Privacy and Autonomy: A Reappraisal," *Southern Journal of Philosophy*. 40(4): 587–604
- Thomson, Judith Jarvis. 1975. "The Right to Privacy". *Philosophy and Public Affairs*. 4(4): 295-314.
- Tene, Omer. Jules Polonetsky. 2013. "Judged by the Tin Man: Individual Rights in the Age of Big Data." *Journal on Telecommunications and High Technology Law*. 11(351).
- Timan, T. B. Newell. B. Koops. 2017. *Privacy in Public Space: Conceptual and Regulatory Challenges*. Cheltenham: Edward Elgar Publishing.
- Turkson, R. E. E. Y. Baagyere and G. E. Wenya. 2016. "A machine learning approach for predicting bank credit worthiness." *Third International Conference on Artificial Intelligence and Pattern Recognition (AIPR)* 1-7.
- Wachter, Sandra. 2019. "Data Protection in the age of Big Data." *Nature Electronics* (2): 6–7.
- Wachter, Sandra. Brent Mittelstadt. 2019. "A Right to Reasonable Inferences: Re-thinking Data Protection Law in the Age of Big Data and AI". *Columbia Law Review* (2): 1-182.
- Wrenn, C. 2007. "Why There Are No Epistemic Duties." *Dialogue*. 46(1): 115-136.

BIG DATA ANALYTICS AND HOW TO BUY AN ELECTION

Jørn Sønderholm, Jakob Mainz, and Rasmus Uhrenfeldt

We show how to lawfully buy an election. The key things that make it possible to buy an election are the existence of public voter registration lists and the existence of Big Data Analytics that can predict how a given elector will vote in an election. Someone interested in buying an election can enter an employment contract with some of the opponent electors where these electors are paid to do a job that prevents them from voting. By purchasing access to public voter registration lists, it is possible to verify *ex post* whether the opponent electors have abstained. In the last two sections, we discuss several barriers that can undermine an attempt to buy an election in the manner we identify.

I. INTRODUCTION

In this article, we show how it is possible to lawfully buy an election.¹ The method we describe for buying an election is novel. The key things that make it possible to buy an election are (1) the existence of public voter registration lists where one can see whether a given elector (E) has voted in a particular election, and (2) the existence of Big Data Analytics (BDA) that with a high degree of accuracy can predict how a given elector will vote in an upcoming election.² Someone interested in buying an election can enter an employment contract with all or some of the opponent electors where these electors are paid to do a job that prevents them from voting.³ By purchasing access to the public voter registration lists, it is possible to verify *ex post* whether the opponent electors that one has signed a contract with have abstained. The method we describe for buying an election is one that revolves around the practice of negative vote buying. “Negative vote buying,” an established term in the literature,⁴ denotes the practice of paying electors to abstain from voting.⁵

The method for buying an election can be used in all electoral districts that have voter registration lists containing information about whether a given elector voted in the latest election and where access to this information is open to the relevant individuals or organizations.⁶ In this article, we describe how the general method for buying an election can be used in a US context. This means that we primarily use empirical data from US elections and US legislation.

Is it a problem that it is possible to buy an election? Some theorists, such as Michael Sandel,⁷ Debra Satz,⁸ and Robert Dahl,⁹ think that money influences politics too much, and it is clear that when one buys an election, money influences politics. They think that money should play a minimal role when it comes to deciding political elections, and they are in favor of rather strict campaign finance regulations.¹⁰

Other theorists are more comfortable with money playing a role in deciding elections and have defended markets in votes. Freiman defends a legal right to buy and sell votes.¹¹ Taylor¹² and Brennan¹³ defend a moral right to buy and sell votes.¹⁴ In this article, we are agnostic about the normative issue of what role money should play in politics in general as well as what role it should play in deciding elections.

The article is structured as follows: In section 2, we give a brief account of voter registration lists and explain what type of information they commonly contain. We also explain how BDA can be used to accurately predict how electors will vote in an upcoming election. In section 3, we offer an example of a generic employment contract that can be used by an individual who is interested in buying an election. In section 4, we discuss the novelty of how the combination of BDA and voter registration lists enables an individual to buy an election. In sections 5 and 6, we discuss several barriers that can undermine an attempt to buy an election in the manner we identify.

Let us end this introductory section by emphasizing that we do not endorse anyone's attempt to buy an election in the manner we describe. With this article, we hope to draw attention to the fact that state-of-the-art BDA makes it possible to buy an election in electoral districts in which there are voter registration lists that contain the voting history of electors. If one finds this possibility disturbing, then there is a *pro tanto* reason to work for the implementation of one, or more, of the three regulatory policy proposals we describe in the final section of the article.

2. VOTER REGISTRATION LISTS AND BIG DATA ANALYTICS

In this section, we explain how voter registration lists work in a US context, and we give an account of the basic features of BDA. In the United States, E has to register with the state she lives in to be able to vote. That is, E has to do something active to get on her state's voter registration list. Which pieces of information are included on the list will vary from state to state. A list can contain information

about E's residential address, gender, registration date, and date of birth.¹⁵ In some states, such as Alaska, the voter registration list does not include information about a voter's date of birth, social security number, voter ID number, place of birth, or signature.¹⁶ In other states, such as Alabama, the only thing the list may explicitly *not* contain is a voter's social security number.¹⁷ In every state, the list contains the voting history of all electors in that state.¹⁸

That voter registration lists contain the voting history of electors means that information about *whether* individual electors have voted in elections is publicly available for a fee. It does not mean that information about *how* individual electors have voted in elections is available. If such information were available, the institution of the secret ballot would be annulled. Annulling this institution would be highly controversial given that it is an entrenched democratic institution.¹⁹

It varies from state to state who may purchase access to the state voter registration list. In New Hampshire, for example, neither researchers nor nonprofit organizations are allowed access, while political committees and candidates are allowed access.²⁰ In Michigan, everyone may purchase access, but only for non-commercial purposes.²¹ A political party can purchase all the information contained in all the voter registration lists for all federal states.²² It will cost an estimated US\$ 136,671 if a political party wants to purchase access to *all* available information from *all* states.²³

Moving on to BDA, it is important to describe some basic features of this type of analytics.²⁴ Using BDA makes it possible to find statistical correlations in big data sets and to make predictions based on these correlations. This means that when an individual (or a political campaign or a company) has access to large amounts of data about electors, and she has state-of-the-art BDA technology at her disposal, she can predict with a high degree of accuracy whether E will vote in an upcoming election and, on the assumption that E will vote, whom she will vote for (which party or candidate). These predictions are often based on widely available data about individual electors that these individuals voluntarily share. These data include data about age, gender, race, income, education level, religious observance, postal address, and what type of car one owns. Owning a pickup truck correlates, for example, with voting for the Republicans, while owning a sedan correlates with voting for the Democrats.²⁵

With access to only twelve data points, it is possible to predict with accuracies exceeding 90 percent for certain demographic groups, whom members of a group will vote for in an upcoming election.²⁶ Note that things get more complicated if we move from a *de facto* two-party system like the one in the United States to a multi-party system like the one in Germany or Israel. In such a multi-party system, a predictive algorithm using only twelve demographic data points is likely to be less accurate than a similar one in a two-party system.

Political parties have access to large amounts of data about electors. For instance, the Republican National Committee and the Democratic National

Committee each have more than nine hundred data points on every American elector.²⁷ With access to digital behavioral data, such as data about whom E is friends with on Facebook or which posts E likes on Facebook, one can get a rather fine-grained picture of the electors' political preferences.²⁸ If political campaigns and/or political parties do not have access to electors' Facebook data (or data from other social media platforms), they can, and often do, hire a private company to collect and analyze these data for them.²⁹

It is important to note that with state-of-the-art BDA, it is not only possible to make accurate predictions of voting behavior at the group level; it is also possible to predict with significant accuracy how individual electors will vote. In an important article in the *Journal of Economic Perspectives*, Nickerson and Rogers describe how political campaigns can use BDA to predict voting behavior at the individual level.³⁰ The BDA methods have only become more sophisticated since 2014, and the method for buying an election we describe in this paper is therefore likely to be even more effective than it would have been in 2014.

3. HOW TO BUY AN ELECTION

To exemplify how it is possible to buy a US election, consider the following example. In a hypothetical electoral district, there are one thousand electors.³¹ They can choose between two candidates from two different parties: candidate D from party ALPHA and candidate R from party BETA. Individual K wants candidate R to win the election. K contacts a data analytics company and asks it to scrape publicly available data, including data generated by social media activity, on each of the one thousand electors. K then asks the company to give a prediction concerning whom each elector will vote for in the upcoming election. She then contacts each elector who is identified as likely to vote for candidate D and asks each such elector if that person is interested in entering an employment contract with her. Here is the generic version of the employment contract K sends out (henceforth, the Employment Contract):

Employment Contract

I, [Name], must on date X, between time Y and time Z, be out of the county in which I officially reside, and I must, throughout the day, be engaged in trash collection in public spaces.

I, [Name], will be eligible for payment (W dollars) as soon as I have signed this contract.³²

In the United States, E can vote in one of three ways. She can vote in person on Election Day, she can vote early by mail, or she can vote early in person. The Employment Contract bars E from voting on Election Day by keeping E away

from her polling station. It does not bar E from voting through the two means of early voting. It is helpful to say more about the Employment Contract.

First, the specific contracts sent out to D electors do not contain the variables X, Y, Z, and W. These variables will be replaced with specific numbers. X is the date of the election. Y is the time at which polling stations in E's home county open. Z is the closing time of polling stations in E's home county. W is the amount of money K offers to E. The contract will be sent to E shortly after the deadline has passed for registration for early voting by mail and early voting in person.

Second, in a US legal setting, K cannot lawfully pay E or offer her any other expenditure not to vote. The reason for this is that it is illegal to offer someone an expenditure in exchange for abstaining. It is also illegal to accept such an offer. Consider this federal law:

Whoever makes or offers to make an expenditure to any person, either to vote or withhold his vote, or to vote for or against any candidate; and Whoever solicits, accepts, or receives any such expenditure in consideration of his vote or the withholding of his vote—Shall be fined under this title or imprisoned not more than one year, or both; and if the violation was willful, shall be fined under this title or imprisoned not more than two years, or both.³³

Prima facie, K does not violate this law by offering the Employment Contract to E since K does not offer E money to withhold her vote. In the Employment Contract, the words “voting” and “vote” do not occur. However, it is an implication of signing the Employment Contract that E cannot vote on Election Day without breaching an employment contract she voluntarily entered.

Another important US federal election law is 42 U.S.C. § 1973i(c). The relevant aspect of it reads like this:

Whoever knowingly or willfully gives false information as to his name, address or period of residence in the voting district for the purpose of establishing his eligibility to register or vote, or conspires with another individual for the purpose of encouraging his false registration to vote or illegal voting, or pays or offers to pay or accepts payment either for registration to vote or for voting shall be fined not more than \$10,000 or imprisoned not more than five years, or both.³⁴

Prima facie, K does not violate this law either. The reason for this is that the law is silent on the issue of whether someone can pay an elector for abstaining. In section 5, we return to the important legal issue of whether K complies with the law. We do this by undertaking a lengthy examination of whether it is likely that US courts will deem the Employment Contract invalid.

Third, the contract E receives from K is delivered electronically, and E can sign it electronically. If E does not sign the contract within 12 hours of receipt, a new contract will be sent to E where the amount of money offered to E has increased. The amount offered to E in the original contract, as well as the increased amount

in the second contract, is something that is determined on an individual basis by the prediction algorithm like the one used to identify E as a likely D-elector. If E does not sign the second contract within the 12-hour window, a third contract will be sent to her with an increased monetary offer. This procedure will continue, with a pre-decided price-ceiling, until the point where either E signs the contract or voting starts in E's home county.

It is important to be aware that for K to be successful in buying the election, she does not need to sign a contract with every D-elector. She only needs to sign a contract with enough D-electors.³⁵ Moreover, K can be successful even if some of the electors who sign the Employment Contract end up voting. The reason for this is that what K needs is only that enough D-electors abstain. How many D-electors are "enough" is relative to each election and depends on a number of variables. We will discuss the most important of these variables in detail later.

Fourth, the company that K has engaged is likely to be able to predict how much money individual D-electors are likely to demand to sign the contract with K. It is this predicted amount that will figure in the original contract that E receives. If the algorithm predicts that E's price is above K's price-ceiling, K offers the maximum amount that she is willing to spend on an individual contract. It is important to stress that the Employment Contract is constructed in such a manner that the cost K incurs from E signing the contract is relatively modest. By signing the Employment Contract, E incurs two types of costs: transportation costs associated with out-of-county travel and opportunity costs. The latter comes in at least two varieties: costs associated with the inability to engage in other paid work on Election Day, and costs associated with the inability to engage in non-paid (social) activities. These costs must be borne by K and must be reflected in the amount of money K offers E to sign the Employment Contract.³⁶ Given that these costs vary from voter to voter, voters whose costs are low are often more attractive to K than voters whose costs are high. Therefore, it is often strategically wise for K to focus her attention on electors within the former group. Also, it would be strategically smart for K not to offer the Employment Contract to D-electors in electoral districts where her favored candidate is likely to win without her interference.

Fifth, it should be recognized that in real-life examples involving K and millions of electors, K's attempt to buy the election will cost a significant amount of money. Each of the contracts K signs may not be overly expensive to K, but the overall amount of money that K must pay for all the contracts she must sign can be vast. How much money K, in the end, must pay to buy the election is to a large extent dependent on empirical circumstances. These circumstances include (i) the number of electors, (ii) the electors' general political sympathies regarding individual candidates/parties, (iii) the strength of these sympathies, (iv) the level of E's transportation and opportunity costs, and (v) the closeness of the election. For example, in an election involving millions of electors, it may be that K only has to sign a few hundred contracts to achieve her desired political outcome. This

will be so when the election is close. On the other hand, in an election involving 100,000 voters, K may have to sign, say, eighty thousand contracts because most of the electors are committed D-electors. An important lesson to be learned from this is that it can be relatively cheap for K to buy an election even if it involves a large electorate, and that it can be relatively expensive for K to buy an election even if it involves a small electorate.

We suggest that for elections that are either close or are ones in which the majority of electors share K's political preferences but a significant part of this majority is likely to abstain, K's method for buying the election is unlikely to require an amount of money that lies beyond the financial resources of some wealthy individuals. In the 2000 US presidential race, George Bush won the state of Florida. He received 537 votes more than Al Gore. Six million votes were cast. This means that Gore would have won the state if 538 Bush voters had abstained on Election Day.³⁷ The point here is that the method we describe for buying an election is empirically feasible in the world we live in, and not only in the theoretical realm. History is full of examples of people being willing and able to pay enormous amounts of money to ensure certain political outcomes. A recent example is Michael Bloomberg's attempt to secure the Democratic Party nomination for the 2020 US presidential election.³⁸ Imagine how much money K would have at her disposal if she joined forces in a consortium with other wealthy R-supporters. If individuals as wealthy as Bloomberg, Forbes, and so on all decided to offer likely D-electors to sign the Employment Contract, the chance of R winning the election would be much higher than in a scenario in which K acts alone. Furthermore, imagine a scenario in which the consortium also included party BETA. This party could then throw a significant part of its campaign resources into the attempt to buy the election.³⁹ This would further amplify R's chance of winning the election.

4. WHAT IS NOVEL ABOUT THIS WAY OF BUYING AN ELECTION?

In this section, we discuss the novelty of how the combination of BDA, voter registration lists, and the Employment Contract enables an individual/consortium to buy an election. The phenomenon of negative vote buying existed long before the advent of BDA. For example, the Aboriginal peoples in Australia were not, as opposed to electors from other demographic groups, mandated to vote from 1962 to 1984. In that period, alcohol was often used to lure away Aboriginal people from the polls.⁴⁰ Between 2012 and 2016, vote buying occurred in 52.2 percent of all elections in Asia, 65.5 percent in post-Soviet countries, and 52.2 percent in Sub-Saharan Africa.⁴¹ However, the recent emergence of BDA makes it possible—on a large scale—to effectively pinpoint *who* the opponent voters are, *how* likely they are to vote for the opponent party, and *how much* money they are likely to demand to abstain. In combination, these three things mean that K

can now buy an election more effectively than before the emergence of BDA.⁴² Big data analytics technology and predictive algorithms are still in their infancy. We conjecture that predictive algorithms will increase in accuracy and that they will increase in accuracy in proportion to how many data points are being added for individual analysis. That is, in the future, predictive algorithms will likely be very accurate in predicting the future voting behavior of electors if they are fed not merely a few data points, but thousands—or even millions of data points.⁴³ Perhaps there will be some diminishing marginal accuracy since, at some point, each new data point begins to correlate so highly with existing data points that each new data point adds less extra accuracy. However, this does not alter the main point here, namely, that the Ks of the future will have at their disposal an improved BDA technology that will make their endeavor to buy an election more effective than K's current endeavor.

It should be noted that BDA has recently been used in an effort to influence voter behavior. The Trump campaign admitted that it ran three voter-suppression campaigns leading up to the 2016 presidential election. Based on huge amounts of personal data, the campaign tried to predict which electors were likely to vote for Clinton, and then encouraged these electors to abstain from voting.⁴⁴ This method is not identical to the method we describe for buying an election, since Trump's voter-suppression campaigns did not involve employment contracts. The method we describe for buying an election is more effective than the method used by the Trump campaign. This is so because K's scheme gives E a stronger incentive to abstain than the one offered by the Trump method. E's incentive is stronger for at least two reasons.

First, note that E accepts the offer of W dollars in exchange for doing a job that prevents her from voting. If receiving W dollars were not a sufficiently strong incentive for E to abstain, then presumably E would not have accepted the Employment Contract.⁴⁵ Second, E knows that if she does not comply with the contract, it can have legal repercussions for her. K can take her to court for breach of contract.⁴⁶ This gives E a strong incentive to comply with the contract, and this is an incentive that is absent in the method used by the Trump campaign. Moreover, from the perspective of the individual behind the attempt to influence voter behavior, K's method has an advantage as compared to the one employed by the Trump campaign. Prior to Election Day, K has more certainty about the outcome of her activities than the Trump campaign has at that point in time. Prior to Election Day, K knows the exact number of people who have signed the Employment Contract. This number gives her a detailed, though imperfect, picture of how many of the electors whom she has targeted will abstain. Prior to Election Day, the Trump campaign has no such picture. To get a sense of the effectiveness of the endeavor to influence voter behavior, the campaign has to rely on imperfect polling data about how electors who have been subjected to the voter suppression campaigns will vote.

5. THE COURTS AS A BARRIER TO K'S SUCCESSFUL ATTEMPT TO BUY AN ELECTION

In this section and the next, we discuss a range of barriers that can undermine K's attempt to buy an election. In this section, we focus on a legal barrier: the courts might decide that the Employment Contract is invalid.⁴⁷ If the courts deem the Employment Contract invalid, the proposed method for buying an election will not work. However, there is reason to think that the courts will deem the Employment Contract valid, on the assumption that the courts generally strive to be consistent. We offer three reasons for being confident that the courts will deem the Employment Contract valid. The first reason is that the courts already deem valid employment contracts that imply that the employee cannot vote. Consider, for instance, the hypothetical elector described below.

Susan is an unemployed chef from Houston. She plans to vote in the 2020 US presidential election on November 3. She registers to vote on October 5, and by October 23, she has not applied for a ballot by mail. Therefore, she cannot vote early by mail. On October 30, she receives an employment contract to work for ExxonMobil on an oil rig. Susan immediately signs the contract, although the deadline has passed for early voting in person. Susan leaves for the Gulf of Mexico on Monday, November 2, and she will be on the oil rig on Election Day. Under Texas law, it is an implication of Susan signing the contract and adhering to its terms that she cannot vote in the 2020 US presidential election.⁴⁸

If the Employment Contract is one that the courts deem invalid, then Susan's employment contract must be a contract that the courts deem invalid. After all, Susan's contract is also one that prevents Susan (the employee) from voting. If we are correct that the courts will not deem Susan's employment contract invalid, then it is, from a perspective of consistency, difficult to see why the courts will not deem the Employment Contract valid. The two employment contracts are identical when it comes to the following two key features: First, in each employment contract, an employer makes an expenditure to an employee. Second, it is an implication of each contract that any employee who signs it cannot vote. The two employment contracts are not identical in all aspects. The type of work that they require from the respective employees differs significantly. It is, however, difficult to see why this difference should be a concern for the courts given that (i) the work required by each of the two contracts is legally permissible, (ii) the work required by each of the two contracts is socially valuable, and (iii) all parties to the contracts are consenting adults.

The second reason why the courts will deem the Employment Contract valid is that it does not violate § 1973i(c), which is one of two relevant federal statutes of US election law. Recall that § 1973i(c) states: "Whoever knowingly or willfully . . . pays or offers to pay or accepts payment either for registration to vote or for voting shall be fined not more than \$10,000 or imprisoned not more than

five years, or both.” In *United States v. Garcia*,⁴⁹ the court upheld a decision from a lower court in which several defendants were convicted of violating § 1973i(c) because they offered welfare food vouchers to voters in return for their promises to vote absentee for certain local candidates in a Democratic primary election.⁵⁰ This court ruling does not apply to the Employment Contract because § 1973i(c) is asymmetrical in the sense that it only outlaws paying someone for voting or registering to vote. It does not outlaw paying someone to abstain from voting. This is important given that what K is doing cannot be interpreted as paying E for voting or registering to vote. Note that the defendants in *United States v. Garcia* were convicted for explicitly offering money, or something of monetary value, to buy votes. This is not what K does. She offers an employment contract. By offering the Employment Contract, K certainly adheres to the letter of § 1973i(c). K would be violating § 1973i(c) if she promised something of value to E in return for E’s promise to vote. K would also be violating § 1973i(c) if she gave something of value to E in exchange for E providing proof that she voted.⁵¹

The third reason why the courts will deem the Employment Contract valid is that it does not violate 18 U.S.C. § 597 (Expenditures to influence voting), which is the second of the two relevant federal statutes of US election law. Recall that § 597, *pace* § 1973i(c), implies that there is legal symmetry between “making an expenditure to make someone vote” and “making an expenditure to make someone withhold her vote.” Now consider election festivals as a get-out-the-vote tactic. Imagine that K uses this tactic. She then proceeds in the following manner: well before Election Day, K selects a voting site that is suitable for an election festival and seeks permission to hold a festival there. She then advertises the upcoming festival at the selected location. The festival essentially features free food and drinks as well as entertainment. The festival is open to everyone, regardless of whether they vote or are eligible to vote. Holding such a festival is legal under federal law, and there is strong empirical evidence that election festivals increase turnout rates significantly.⁵² It is evident from the literature on compulsory voting that making sure that electors turn up at polling stations is something that almost always make them vote.⁵³ It is therefore not surprising that if holding election festivals is something that attracts electors to the physical vicinity of polling stations, holding such festivals increases turnout. Green and McClellan also gauge the cost-per-vote of election festivals. On certain assumptions about the cost of the election festival as well as the number of registered voters in the electoral district in which the festival takes place, “festivals generate approximately 43.7 votes per precinct at \$48 per vote, which is quite good by the standards of rigorously evaluated get-out-the-vote programs.”⁵⁴

In a legal environment in which there is symmetry between making an expenditure to make someone vote and making an expenditure to make someone withhold her vote, and in which election festivals are legal, it is reasonable to

suppose that the courts will deem the Employment Contract valid. What are the potential disanalogies between holding an election festival and doing what K does in terms of sending out the Employment Contract? (1) If K holds an election festival, she does not make an expenditure to anyone in return of proof of vote. By sending out the Employment Contract, K does not do this either. (2) If K holds an election festival, she announces the festival before Election Day. K also sends out the Employment Contract before Election Day. (3) If K holds an election festival, she will likely make an expenditure to individuals who do *not* vote. If K sends out the employment contract, she will likely make an expenditure to individuals who *do* vote. *Prima facie*, this might seem like a disanalogy between the two activities. Here, it is, however, of crucial importance to remember that the two activities are being evaluated with reference to § 597, which treats making an expenditure to make someone vote symmetrically to making an expenditure to make someone withhold her vote. (4) If K holds an election festival in a voting district with historically low turnout and in which the vast majority of electors have a particular political preference (K can buy access to both types of information), K is willfully engaged in an activity that she knows is likely to have a partisan bias (favor one candidate/party over other candidates/parties). If K sends out the Employment Contract to electors with a particular profile (as she does), K is willfully engaged in an activity that she knows is likely to have a partisan bias (favor one candidate/party over other candidates/parties). (5) If K holds an election festival, K is likely to do this with the intention of increasing turnout among a select group of electors. If K sends out the Employment Contract, K is likely to do this with the intention of decreasing turnout among a select group of electors. (6) If K holds an election festival, K will be making an expenditure to identifiable electors (those electors who receive the free food/beverages or enjoy the entertainment at the festival). If K sends out the Employment Contract, K will be making an expenditure to identifiable electors (those electors who sign the contract).⁵⁵

Of course, if K sends out the Employment Contract and uses state-of-the-art BDA to decide whom to send the contract to, she can affect turnout and thereby increase the chances that her preferred political result will materialize much more effectively than if she organizes an election festival. It is, however, difficult to see how this disanalogy between the two types of activities could be legally relevant in light of § 597 and § 1973i(c).

Taking these arguments about the similarity between the Employment Contract and Susan's employment contract, the restricted scope of § 1973i(c) and the analogies between holding election festivals and sending out the Employment Contract in a legal environment in which § 597 is in place make it reasonable to proceed on the assumption that the courts will deem the Employment Contract valid.

6. OTHER BARRIERS TO K'S SUCCESSFUL ATTEMPT TO BUY AN ELECTION

Another set of barriers that can undermine K's attempt to buy an election revolves around modes of voting. The Employment Contract does not bar E from voting early. This means that all the D-electors who vote early are voters whom K cannot get to abstain by getting them to sign the Employment Contract. A significant number of D-electors are simply out of reach for K. To exemplify this problem, recall the example involving one thousand electors, and candidate D from party ALPHA, and candidate R from party BETA. Suppose that there are six hundred D-electors and four hundred R-electors. The latter will all vote. Assume that 401 of the D-electors vote early. Then, even if K has a perfect success rate and signs a contract with each of the remaining 199 D-electors, candidate D wins. This means that K was not able to buy the election. This general problem for K involving D-electors voting early can be exemplified in countless other ways involving different numbers, and it shows that when a big enough subset of the electorate behaves in a particular manner, K cannot buy an election. It is also worth noting here that K cannot control whether a big enough subset of the electorate behaves in the relevant manner.⁵⁶

Regulators can amplify the "voting-early" barrier by making it easier to vote early. Making early voting easier can be done in at least two ways. The period in which such voting is possible could be extended. Moreover, current US rules to the effect that one has to sign up/apply for permission to vote early could be scrapped such that by being eligible to vote automatically makes one eligible to vote early.

Consider next regulations for voting during working hours. In several US states, there are laws in place that give employees the right to take time off to vote. Such a right potentially undermines K's attempt to buy an election. Consider, for example, the scenario in which E signs the Employment Contract, and, on Election Day, takes time off to vote for D. In this scenario, K has paid E, E has not abstained, and K has no grounds for legal complaint against E. If enough other D-electors who have signed the Employment Contract also exercise their right to take time off to vote, then K's attempt to buy the election fails. How big a problem for K is it that electors in some states have the right to take time off to vote? Note that as of 2020, there are nineteen states in which employees are *not* entitled to take time off to vote.⁵⁷ In these states, K's attempt to buy a state election cannot fail because of employees' right to take time off to vote. Moreover, this aspect of US employment contract law also leaves intact K's attempt to buy a nationwide election. It might be that what K must do in order to buy such an election is to focus her efforts on only those nineteen states where electors do not have the right to take time off to vote. Importantly, a subset of these nineteen states consists of swing states.

Further, note the distinction between having a right to X and exercising the right to X. This distinction is important because what potentially undermines K's attempt to buy the election is the exercise of the right to take time off to vote—and not merely having this right. So, even in the states in which employees who have signed the Employment Contract have the right to take time off to vote, it is possible that these employees (or many of them) do not exercise their right. If none (or few) of them do not, K does not have a problem. The question of how many of the electors who have signed the Employment Contract will exercise their right to take time off to vote is an empirical one, which we are not in a position to answer.

It is, however, likely that not everyone who has the right will exercise it. The states where employees have a right to take time off to vote are of two types. There are states in which the employer pays for the time that the employee takes off to vote, and there are states in which the employee pays for this herself. In the latter states, of which there are seven, there is a significant opportunity cost associated with exercising the right to take time off to vote. This cost consists of a loss of personal income. The higher W is, the higher the amount is that E loses from exercising her right, and the more likely it is that E will not exercise her right.⁵⁸

Regulators can amplify the “voting-during-working-hours” barrier by changing US contract law such that all states give employees the right to take time off to vote. The issue of whether it is the employer or the employee who should pay for the time taken off to vote is a secondary one. For reasons laid out above, it is, however, likely that if this cost is borne by the employer, then more employees will exercise their right than if this cost is borne by the employee.

A further barrier to K's successful attempt to buy an election is that E can sign the Employment Contract and then claim to be sick on Election Day.⁵⁹ However, if E calls in sick, then E is supposed to stay at home. Now, if E calls in sick and goes voting, K will know about E going voting by checking the voter registration lists, and K can approach E afterward and say: “Dear E, you called in sick on Election Day and then went voting. If you are too sick to work, how come you were not too sick to go voting? I will now take legal means to try to recover the money I paid you. I feel that you have breached the contract.” Of course, K might not want to jump through all these hoops, but the important fact is that E can know that K has access to her voting records and can check whether E voted. If E has this knowledge, then this is something that is likely to make her hesitant to call in sick and then go voting.

Another barrier is that the polls are typically open for a longer period of time than the daily hours of a regular job. It might be illegal for K to offer an employment contract with such long hours that complying with the contract prevents the employee from voting.⁶⁰ However, the Fair Labor Standards Act (FLSA) of 1938 dictates policy for most workers. According to an interpretation of the FLSA

by the US Department of Labor, the act does not limit the number of hours in a day or days in a week an employee must work, including overtime hours, if the employee is at least 16 years old.⁶¹

Let us end this article by describing a third way in which regulators can undermine K's attempt to buy an election. Regulators can eliminate the voting history of individual electors from the voter registration lists. It is information about whether E has voted in the latest election that makes it possible for K to *ex post* verify whether E has fulfilled her contractual obligations. If E voted, then K knows that E violated the Employment Contract, and K can take legal action against E. If E did not vote, then K does not know if E has fulfilled her contractual obligations. After all, not voting is compatible with staying at home on Election Day and not collecting any trash in a county other than one's home county. However, K is likely to have no objections to this way of breaching the Employment Contract given that K's only interest is that E abstains.

If K has no way of verifying whether E has fulfilled her contractual obligations, entering the Employment Contract with E becomes nothing more than a gamble for K. There are two ways of eliminating the voting history of individual electors from the voter registration lists. The first is to eliminate the voter registration lists themselves. The second way is to keep the voter registration lists, but reform them such that the voting history of individual electors is removed from them.

We acknowledge that in addition to checking the voter registration lists, there are at least two ways in which K can verify that E has fulfilled her contractual obligations. First, K can dispatch a supervisor to monitor E's work on Election Day. This is not an attractive option for K given that she has to send out a multitude of supervisors on Election Day to a multitude of locations. This will greatly increase K's overall costs. Also note that the type of work involved in the Employment Contract is such that K cannot verify that the work has been done by checking on the day after Election Day. K cannot verify whether E has been collecting trash in the relevant time frame because K cannot measure how much trash there was at the beginning of the workday and then hold this up against her estimate of how much trash is lying around the day after. Second, E can include in the contract a requirement that E must install a tracing application on her phone and have it turned on throughout her working shift. K can then verify the whereabouts of E on Election Day and thereby verify that E has been out of her home county and away from her polling station. The use of tracing applications will not allow K to verify that E has been collecting trash, but given that the whereabouts of E on Election Day are what is important to K, this issue is of no real importance to K.

There are at least three reasons as to why the use of tracing applications is not an attractive option for K. First, K has to spend money on developing a tracing application of her own, and there are costs associated with checking the data generated by all the electors who signed the Employment Contract. Second, it is likely to significantly increase E's psychological discomfort associated with complying

with the Employment Contract if she has to consent to use a tracing application on Election Day. Given that K has to compensate E for this discomfort, the use of tracing applications increases K's overall costs. Put briefly, E is likely to demand something in return (more money) for giving K access to detailed information about her physical location. Third, it is easy for E to game K's use of a tracing application to verify whether E has fulfilled the Employment Contract. E can get somebody to take her phone and drive out of E's home county. E can also drive there herself, leave her phone somewhere (or pay someone to walk around with it for the whole day), drive back (and vote) and then collect her phone after the working shift. K can, of course, close these loopholes by requiring that the type of tracing application that E uses must be an ankle bracelet monitor that only K can fit and remove from E's foot. We conjecture that such a requirement hugely increases E's psychological discomfort, and that many electors, as a result of this increased discomfort, will either not sign the Employment Contract or demand a lot of money to do so. Such a demand drives up K's costs.

It is clear from the discussion in the preceding two sections that there are several conditions under which K cannot buy the election. For example, the courts deem the Employment Contract invalid; too many D-electors vote early; or too many D-electors, who sign the Employment Contract, take time off to vote on Election Day. Therefore, the method for buying an election described in this article is not foolproof. It does not always work. This fact does not, however, render the article irrelevant or uninteresting. It is still the case that there are many conditions under which K can buy an election and, importantly, some of these conditions obtain in real-life elections.

Aalborg University

NOTES

We wish to thank Sarah Birch, Richard Hasen, Alexandru Volacu, Jason Brennan, and an anonymous reviewer from *Public Affairs Quarterly* for constructive and helpful comments on earlier versions of this paper.

1. To clarify, for x to “buy an election” means that x makes an expenditure to individual electors to influence their voting behavior (either to vote *simpliciter*, to abstain, or to vote for a particular candidate/party/proposition) such that x 's preferred election result is brought about.

2. In this article, we use “elector” to denote a person who has the right to vote in an election. The term is not used to denote a member of the US Electoral College.

3. For any individual/entity x , an “opponent elector” is any elector that votes, or intends to vote, for a candidate/party/proposition that is different from the one x supports.

4. Morgan and Várdy, “Negative Vote Buying.”

5. In the United States, there are mobile apps that allow you to see whether individual electors have voted. See, for example, the app VoteWithMe at <https://bit.ly/2x3Gfll>, or the app Outvote at <https://bit.ly/2IXEzfQ> (both accessed November 9, 2020).

6. This is on the assumption that BDA is available in these electoral districts. Moreover, as we will explain in detail later, there are several conditions under which this method can be used, but will be unsuccessful.

7. Sandel, *What Money Cannot Buy*.

8. Satz, *Why Some Things*.

9. Dahl, *On Democracy*.

10. For an illuminating discussion of various methods for limiting the effect of money on politics (including the effect of limiting private contributions to candidates or parties), see Christiano (“Money in Politics”).

11. Freiman, “Vote Markets.”

12. Taylor, “Two (Weak) Cheers.”

13. Brennan, *Ethics of Voting*.

14. Moreover, Volacu argues that there are at least two plausible *prima facie* reasons in favor of barter voting markets (“Electoral *Quid Pro Quo*”). Note, however, that in a barter voting market, money is not exchanged, and therefore money cannot be said to influence politics or decide elections.

15. See https://sos.ga.gov/index.php/elections/order_voter_registration_lists_and_files (accessed November 9, 2020).

16. See <http://voterlist.electproject.org/states/alaska> (accessed November 9, 2020).

17. See <http://voterlist.electproject.org/states/alabama> (accessed November 9, 2020).

18. See <http://voterlist.electproject.org/states>. See also <https://www.ncsbe.gov/results-data/voter-history-data> (accessed November 9, 2020).

19. Birch and Watt, “Remote Electronic Voting,” 62; Dahl, *On Democracy*, 96

20. See <http://voterlist.electproject.org/states/new-hampshire> (accessed November 9, 2020).

21. See <http://voterlist.electproject.org/states/michigan> (accessed November 9, 2020).

22. See <https://bit.ly/2ThiLk4> (accessed November 9, 2020).

23. See <http://voterlist.electproject.org/full-list-purchase-facts-and-info> (accessed November 9, 2020).

24. A recent definition in the literature regards BDA as “a new generation of technologies and architectures, designed to economically extract value from very large volumes of a wide variety of data, by enabling high velocity capture, discovery and/or analysis.” See Mikalef et al. (“Big Data Analytics Capabilities,” 273).

25. See Gebru et al. (“Using Deep Learning”).

26. See, for example, the tool on the *New York Times*’s website, where you can plot in demographic data about an elector, and the tool will tell you how likely the elector is to

vote for a certain party: <https://nyti.ms/2HHc8lN> (accessed November 9, 2020). See also a similar tool on *The Economist's* website at <https://econ.st/3c19aqr> (accessed November 9, 2020).

27. Moore, "Protecting Democratic Legitimacy."

28. Kosinski et al., "Personality and Website Choice."

29. Susser, Roesller, and Nissenbaum, "Online Manipulation."

30. Nickerson and Rogers, "Political Campaigns and Big Data."

31. The average size of a US congressional district, based on the 2010 Census apportionment population, is 710,767. See <https://www.census.gov/prod/cen2010/briefs/c2010br-08.pdf> (accessed December 13, 2020).

32. The exact type of work that the Employment Contract requires is not of crucial importance to the general point we make in this article. The Employment Contract could, for example, require that E sit at home all day on Election Day making telephone calls asking for donations to K's charity fund that pays for pediatric cancer treatment. Such an employment contract could even allow E to go to work on Election Day and merely require that E goes straight to work and goes straight home at the end of the workday.

33. 18 U.S. Code, § 597. See <https://www.govinfo.gov/app/details/USCODE-2010-title18/USCODE-2010-title18-partI-chap29-sec597> (accessed December 15, 2020).

34. 42 U.S.C. § 1973i(c). See <https://www.govinfo.gov/content/pkg/USCODE-2008-title42/html/USCODE-2008-title42-chap20-subchapI-A-sec1973i.htm> (accessed December 15, 2020).

35. K has access to the complete voting history of E. This information is useful for K, given that she tries to influence E's voting behavior. For example, if K's algorithm predicts that 60-year-old E is likely to be a committed D-elect, but is given the input, from the voting registration list, that E has never voted, then the algorithm can be designed to ignore E (bar extraordinary circumstances such as E being explicit in social media posts that she *will* vote in the upcoming election).

36. How much does it cost to buy a single vote? This is a difficult question. The answer depends on what country and what economic context the attempted vote buying occurs in. In Chonburi province in Thailand, the going rate is reported to be around US\$ 9 per vote, but can rise as high as US\$ 90 (Cheeseman and Klaas, *How to Rig an Election*, 65). The reported prosecutions for US vote buying suggest that minimal payments are involved: US\$ 3 or US\$ 5, in one case. A US\$ 45 welfare voucher and a six-pack in another case, and US\$ 20 and US\$ 30 in other transactions (Karlan, "Not by Money," 1459). Hasen reports that before the rise of the secret ballot, costs were higher in constant dollars ("Vote Buying," 1329).

37. For statistics on this election, see <https://www.npr.org/2018/11/12/666812854/the-florida-recount-of-2000-a-nightmare-that-goes-on-haunting?t=1604406223180> (accessed November 9, 2020).

38. Bloomberg's failed US presidential campaign in 2020 cost him over US\$ 500 million. <https://www.businessinsider.sg/things-mike-bloomberg-bought-in-failed-500-million-presidential-campaign-2020-3?r=US&IR=T> (accessed November 9, 2020). Steve Forbes spent more than US\$ 76 million of his own money on an unsuccessful run at the US presidency in 2000—over US\$ 113 million in 2018 dollars. Donald Trump spent US\$ 66.1 million of

his own money on the 2016 US presidential election. Ross Perot spent US\$ 65.4 million of his own money on the 1992 US presidential election, and Rick Scott ran for the US Senate in 2018 and spent US\$ 63.6 million of his own money. <https://www.tampabay.com/florida-politics/buzz/2018/12/10/rick-scott-spent-63-5-million-on-his-u-s-senate-campaign-where-does-that-rank-all-time-among-political-self-funders/> (accessed November 9, 2020).

39. In the 2020 US presidential election, Joe Biden is expected to be the first candidate in American history to raise US\$ 1 billion through his campaign. <https://edition.cnn.com/2020/10/29/politics/2020-election-cost-money-trump-biden/index.html> (accessed November 9, 2020).

40. Morgan and Várdy, “Negative Vote Buying”; Orr, “Dealing in Votes.”

41. Cheeseman and Klaes, *How to Rig an Election*, 250.

42. Big data analytics also make it possible to predict how likely it is that E is persuaded by a given piece of political advertisement (Papakyriakopoulos et al., “Social Media and Microtargeting”).

43. It is likely that the development of predictive algorithms geared toward forecasting voter behavior follows a general development path that mimics that of predictive algorithms geared toward forecasting, for example, future crime spots, which students are likely to drop out, which inmates are likely to re-offend when released, and epidemic outbreaks around the world. Predictive algorithms geared toward forecasting these things have increased significantly in terms of accuracy. Consider, for example, this estimate from within the field of medicine:

Previous generations of algorithms were largely rule-based models, often requiring manual input of usually <10 variables, to provide clinical decision support for specific situations, such as guiding imaging for pulmonary embolism, with reasonable discrimination and calibration. Over the past 5 years, modern AI-based algorithms have enabled automated real-time prediction based on almost unlimited numbers of variables, with predictive performance superior to that of traditional algorithms. (Parikh, Obermeyer, and Navathe, “Regulation of Predictive Analytics,” 810)

44. Moore, “Protecting Democratic Legitimacy,” 96. Note that the Trump campaign did something that K is also doing: namely, they tried to predict how individual electors will vote, and tried to influence the voting behavior of these individual electors.

45. This is on the assumption that E knows that signing the Employment Contract prevents her from voting.

46. This is on the assumption that E knows that K can verify *ex post* whether E voted.

47. United States election law is complex. There are federal laws that govern elections in which one or more federal candidates are on the ballot. However, there are also state laws. These laws govern elections in which only state candidates are on the ballot. In this section, our focus is strictly on federal law. For an overview of the many state laws that govern state elections, see Hasen (“Vote Buying,” 1324n1).

48. Please see the following links for important 2020 dates for voting in Texas: <https://www.sos.state.tx.us/elections/voter/important-election-dates.shtml#2020> (accessed November 9, 2020).

49. *United States v. Garcia*, 719 F.2d 99 (5th Cir. 1983).

50. For the court opinion associated with the court's verdict, see <https://law.justia.com/cases/federal/appellate-courts/ca9/19-10073/19-10073-2020-09-10.html> (accessed December 9, 2020).

51. This point about providing proof of voting is an important one in US federal election law. It is legal for businesses/nonprofit organizations to hand out free food and/or offer discounts on various types of commercial products on Election Day as part of festivals/campaigns celebrating voting and/or democracy. Such businesses/nonprofit organizations must, however, make their offer available to everyone, and they cannot demand proof of voting as a requirement for receiving gifts or discounts. See <https://www.bolderadvocacy.org/wp-content/uploads/2016/04/Can-a-Nonprofit-Provide-Incentives.pdf>; <https://www.wsj.com/livecoverage/election-live-updates-trump-biden-2020-10-30/card/Ag7pZgu79eW5Z5ZMSk28> (accessed December 8, 2020). Some state laws permit activities that are not permitted under federal law. For example, Alaska law does prohibit a person from paying another person to vote for a particular candidate or proposition, but no Alaska statute prohibits a person from compensating another person for voting *per se* (for example, by reimbursing an elector for the cost of the fuel that she used to drive to the poll station (Hasen, "Vote Buying," 1326). Also, the Mississippi Supreme Court upheld the right of a candidate in a local election to hold a cash draw close to the voting precinct. On Election Day, 1,279 voters entered the cash-draw by signing a card and placing it in a box. A person at the box asked each entrant whether she had voted and instructed her to vote prior to placing her card in the box. After the polls closed, the drawing was held and the prize money was distributed to eleven winners. See *Naron v. Prestage*, 469 So. 2d 83 (1985) at <https://law.justia.com/cases/mississippi/supreme-court/1985/56113-0.html> (accessed November 16, 2020).

52. Green and McClellan, "Election Festivals."

53. Brennan and Hill, *Compulsory Voting*, 6; Birch and Watt, "Remote Electronic Voting," 60–72; Louth and Hill, "Compulsory Voting in Australia," 27.

54. Green and McClellan, "Election Festivals," 5.

55. In the eyes of the courts, what K is doing is different from what she would be doing if she merely walked around with a poster saying "Do Not Vote!" or made an expenditure to other people to make them either walk around with such a poster or broadcast the "Do Not Vote!" message on radio, TV, or social media platforms. If K did this, she would not be making an expenditure directly to individual voters to influence their voting behavior.

56. Most electors in a US presidential election do not vote early, though the trend is that more and more electors are voting early. In the 2016 US presidential election, slightly more than 40 percent of all electors voted early. See <https://www.eac.gov/documents/2017/10/17/eavs-deep-dive-early-absentee-and-mail-voting-data-statutory-overview> (accessed November 9, 2020).

57. See <https://www.businessinsider.com/take-time-off-from-work-to-vote-state-guide-2020-10?r=US&IR=T> (accessed November 9, 2020).

58. Recall that "W" denotes the amount of money offered to E for her to sign the Employment Contract.

59. We thank Sarah Birch for bringing this point to our attention.

60. We thank Sarah Birch for bringing this point to our attention.

61. See <https://legalbeagle.com/7736423-many-can-legally-work-day.html> (accessed December 3, 2020).

REFERENCES

- Birch, Sarah. *Full Participation : A Comparative Study of Compulsory Voting*, Tokyo: United Nations University Press, 2009.
- Birch, Sarah, and Bob Watt. "Remote Electronic Voting: Free, Fair and Secret?" *Political Quarterly* 75, no. 1 (2004): 60–72.
- Brennan, Jason. *The Ethics of Voting*. Princeton, NJ: Princeton University Press, 2011.
- Brennan, Jason, and Lisa Hill. *Compulsory Voting: For and Against*. New York: Cambridge University Press, 2014.
- Cheeseman, Nicholas, and Brian P. Klaas. *How to Rig an Election*. New Haven, CT: Yale University Press, 2018.
- Christiano, Thomas D. "Money in Politics." In *The Oxford Handbook of Political Philosophy*, edited by David Estlund, 241–57. Oxford, UK: Oxford University Press, 2012.
- Dahl, Robert A. *On Democracy*. New Haven, CT: Yale University Press, 1998.
- Freiman, Christopher. "Vote Markets." *Australasian Journal of Philosophy* 92, no. 4 (2014): 759–74.
- Gebru, Timnit, Jonathan Krause, Yilun Wang, Duyun Chen, Jia Deng, Erez Lieberman Aiden, and Li Fei-Fei. "Using Deep Learning and Google Street View to Estimate the Demographic Makeup of Neighborhoods across the United States." *Proceedings of the National Academy of Sciences* 114, no. 50 (2017): 13108–13.
- Green, Donald P., and Oliver A. McClellan. "Election Festivals and Voter Turnout: An Overview of Recent Research." March 4, 2020. Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3548959.
- Hasen, Richard L. "Vote Buying." *California Law Review* 88, no. 5 (2000): 1323–71.
- Karlan, Pamela S. "Not by Money but by Virtue Won? Vote Trafficking and the Voting Rights System." *Virginia Law Review* 80, no. 7 (1994): 1455–75.
- Kosinski, Michal, David Stillwell, Pushmeet Kohli, Yoram Bachrach, and Thore Graepel. "Personality and Website Choice." ACM Web Sciences Conference, Evanston, IL, January 2012.
- Louth, Jonathon, and Lisa Hill. "Compulsory Voting in Australia: Turnout with and without It." *Australian Review of Public Affairs* 6, no. 1 (2005): 25–37.
- Mikalef, Patrick, Maria Boura, George Lekakos, and John Krogstie. "Big Data Analytics Capabilities and Innovation: The Mediating Role of Dynamic Capabilities and Moderating Effect of the Environment." *British Journal of Management* 30, no. 2 (2019): 272–98.
- Moore, Martin. "Protecting Democratic Legitimacy in a Digital Age." *Political Quarterly* 90, no. 51 (2019): 92–106.
- Morgan, John, and Felix Várdy. "Negative Vote Buying and the Secret Ballot." *Journal of Law, Economics, and Organization* 28, no. 4 (2010): 818–49.
- Nickerson, David W., and Todd Rogers. "Political Campaigns and Big Data." *Journal of Economic Perspectives* 28, no. 2 (2014): 51–73.

- Orr, Graeme. "Dealing in Votes: Regulating Electoral Bribery." PhD thesis, Griffith Law School, Griffith University, 2005.
- Papakyriakopoulos, Orestis, Simon Hegelich, Morteza Shahrezaye, and Juan Carlos Medina Serrano. "Social Media and Microtargeting: Political Data Processing and the Consequences for Germany." *Big Data & Society* 5, no. 2 (2018): 2053951718811844.
- Parikh, Ravi B., Ziad Obermeyer, and Amol S. Navathe. "Regulation of Predictive Analytics in Medicine." *Science* 363, no. 6429 (2019): 810–12.
- Sandel, Michael. *What Money Cannot Buy: The Moral Limits of Markets*. New York: Farrar, Straus and Giroux, 2013.
- Satz, Debra. *Why Some Things Should Not Be For Sale: The Moral Limits of Markets*. New York: Oxford University Press, 2010.
- Susser, Daniel, Beate Roessler, and Helen Nissenbaum. "Online Manipulation: Hidden Influences in a Digital World." *Georgetown Law Technology Review*, December 23, 2018. Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3306006.
- Taylor, James Stacey. "Two (Weak) Cheers for Markets in Votes." *Philosophia* 46, no. 1 (2018): 223–39.
- Volacu, Alexandru. "Electoral *Quid Pro Quo*: A Defence of Barter Markets in Votes." *Journal of Applied Philosophy* 36, no. 5 (2019): 769–84.

ISSN (online): 2246-123X
ISBN (online): 978-87-7210-968-8

AALBORG UNIVERSITY PRESS