



Provided by the author(s) and University of Galway in accordance with publisher policies. Please cite the published version when available.

Title	A value-centered approach to data privacy decisions
Author(s)	Carter, Sarah E.
Publication Date	2024-02-06
Publisher	NUI Galway
Item record	http://hdl.handle.net/10379/18048

Downloaded 2024-05-09T08:03:52Z

Some rights reserved. For more information, please see the item record link above.





OLLSCOIL NA
GAILLIMHE
UNIVERSITY
OF GALWAY

A Value-Centered Approach to Data Privacy Decisions

Sarah E. Carter, M.Sc.

Student Number: 19240169

A thesis submitted to the School of History and Philosophy, College of Arts, Social Sciences, and Celtic Studies, University of Galway, in fulfillment of the requirements for the degree of Doctor of Philosophy.

Month of Submission: September 2023

Supervisors:

Dr. Heike Felzmann, Prof. Dr. Mathieu d'Aquin, Prof. Dr. Kathryn Cormican, Dr. Dave Lewis

Graduate Research Committee:

Prof. Dr. John Breslin, Dr. Karen Young, Dr. John Danaher

*For my grandfather, the Rev. Dr. John S. Carter, who
completed his Doctor of Ministry degree in 1976*

Table of Contents

<i>Declaration</i>	<i>i</i>
<i>Funding</i>	<i>ii</i>
<i>Acknowledgements</i>	<i>iii</i>
<i>Abstract</i>	<i>vi</i>
<i>Research Outputs Related to this Thesis</i>	<i>viii</i>
<i>List of Abbreviations</i>	<i>xi</i>
<i>List of Tables</i>	<i>xii</i>
<i>List of Figures</i>	<i>xiii</i>
Chapter 1 <i>Thesis Summary and Structure</i>	1
Section 1.1 Chapter Overview	1
Section 1.2 Thesis Summary: Research Motivation, Aims, and Contributions	1
Section 1.3 Structure of Thesis	2
Chapter 2 <i>Background</i>	4
Section 2.1 Chapter Overview	4
Section 2.2 The Breakdown of Privacy Self-Management	5
Section 2.3 Personal Values: From Self-Management to Self-Governance	12
Section 2.4 From Theory to Practice: Designing for Value-Centered Privacy Decisions	15
Section 2.5 Chapter Summary	24
Chapter 3 <i>Crafting the Value-Centered Approach to Privacy Decisions</i>	25
Section 3.1 Chapter Overview	25
Section 3.2 Conceptualizing Value-Centered Privacy Decisions	26
Section 3.3 Designing a Value-Centered Privacy Assistant	37
Section 3.4 Conclusion	45
Chapter 4 <i>Methodology for Empirical Studies</i>	47
Section 4.1 Chapter Overview	47
Section 4.2 Phase I: Online Value and Privacy Preference Survey	49
Section 4.3 Phase 2: Mock App Store Study	53
Section 4.4 Phase 3: Post-Study Interviews	59
Chapter 5 <i>How Do We Value Data Privacy?</i>	64
Section 5.1 Chapter Overview	64
Section 5.2 Online Value and Privacy Preference Survey	66
Section 5.3 Semi-Structured Interviews	75
Section 5.4 Conclusion	116

Chapter 6	<i>Evaluating the Value-Centered Privacy Assistant</i>	118
Section 6.1	Chapter Overview	118
Section 6.2	The VcPA and Value-Centered App Choices	120
Section 6.3	VcPA Profiles	120
Section 6.4	Selective Notices and the “Suggest Alternatives” Feature	124
Section 6.5	Exploratory Notices	127
Section 6.6	Conclusion	128
Chapter 7	<i>Conclusion and Future Directions</i>	129
Section 7.1	Chapter Overview	129
Section 7.2	Overview of Research Findings and Contributions	130
Section 7.3	Implications for Future Research	133
Section 7.4	Concluding Thoughts	140
Bibliography		142
Appendices		152
Appendix I:	Online Value and Privacy Preference Survey	152
Appendix II:	Apps in Mock App Store	158
Appendix III:	Entrance Survey	161
Appendix IV:	Exit Survey	163
Appendix V:	Pre-Interview Survey	164
Appendix VI:	Study Demographics	165
Appendix VII:	Heatmap of Significant Correlations for Loselt! and OpenLitterMap	166
Appendix IIX:	Question Bank for Semi-Structured Interviews	167
Appendix IX:	Additional Quotations from Interviews	168
Appendix X:	Feedback Comments from Exit Survey	173
Appendix XI:	Reported Participant Rationale for Ignoring VcPA Selective Notices	175

Declaration

I, Sarah Elizabeth Carter, confirm that I have not obtained a degree in the University of Galway or elsewhere based on the work contained in this thesis. I am the sole author of this thesis. Any contributions to this work by others have been noted at the beginning of each chapter (under “Collaborator Contributions”).



Signature

September 28th, 2023

Date

Funding

This work was conducted with the financial support of the Science Foundation Ireland Center for Research Training in Digitally-Enhanced Reality (d-real) under Grant No. 18/CRT/6224.

Centre for
Research
Training



The funding body had no role in study design, data collection, analysis, decision to publish, or the preparation of this thesis.

Acknowledgements

*There are places I'll remember
All my life, though some have changed
Some forever, not for better
Some have gone and some remain*

*All these places had their moments
With lovers and friends, I still can recall
Some are dead and some are living
In my life, I've loved them all*

A PhD is a large undertaking, and often a lonely one. I am very grateful to have had such a large international circle of love and support on this journey – across countries, through a pandemic, and out the other side.

Firstly, in my birth country, the United States. Although this thesis is dedicated to one “Dr. Carter,” there are two other “Dr. Carters” in my life I must thank – my dad, Dr. David Carter and my mom, Dr. Wendy (Kauffman) Carter. Thank you both for your inspiring love of learning and steadfast support of my PhD and life abroad. To my sister, Annika, and brother-in-law, Justen – thank you for your love, your support, and your sense of humor. To my grandparents and extended family – thank you for the letters, messages, and calls, and for always embracing me with love when I come home, no matter how long my absence. To the Pragmatic Buddhist meditation groups (virtually in Ohio and Scotland), I thank for providing the space to ground and center during the ups and downs of a PhD. And lastly, to my late dog and virtual conference buddy during the pandemic – little Rusty. You always saw strength in me when no one else did. I love you and miss you every day.

In Spain, *muchos besos a Remi, Juan, Rosa, Emilio, y Pol* for treating me as a member of their family, providing love, support, and (a few!) beach visits during this PhD journey. *Gracias por todo. ¡Os quiero muchísimo!*

In the Netherlands, many thanks to Ilaria, Dayana, and the Knowledge, Representation, and Reasoning (KRR) group for welcoming me to VU Amsterdam as a visiting researcher for a few months. In particular, I would like to thank Mojca for helping with the paperwork and other tasks to make sure I can come to Amsterdam, and KRR members, Romana, Taraneh, Benno, Inés, Tae, Márk, Jan-Christoph, Emil, Michael, Nikos, Dimitris, Lise, Loan, Andreas, and Frank for including me in group activities and really making me feel like part of the team. I would also like to thank Emma, for our in-depth chats, for opening her home to me, and for our wonderful time traveling to FAccT2022 and exploring Seoul together (including dancing “Gangnam Style” in Gangnam). I also am grateful to Gijs (in the Netherlands) and Laura (in Germany) for their feedback during our peer PhD paper review sessions, and Xengie (in Luxemburg) and Marcu (in the Netherlands) for our interdisciplinary discussions and collaboration. I would also like to acknowledge all those who helped me find my true passion at the edge of science and ethics during my master’s, laying the initial foundation that made this PhD possible. Thank you to Mike, Sarah, Karin, Annelien, and the Julius Center at UMC Utrecht for allowing me to intern with them and “do some ethics.” Thank you to Jeff, Eyleen, Marne, and the entire Beekman group for introducing me to Mike and helping me find the research that makes me want to get out of bed in the morning. And, lastly, special thanks to the entire Cancer, Stem Cell, and Developmental Biology (CSDB) master’s community; (now retired) CSDB program leader

Joost; my “Dutch Parents” Pauline and Rocco; the Mount Holyoke Alumnae group in the Netherlands; and friends Esmée, Eva, and Anna for all their support, encouragement, and advice during the journey to find what I love to do.

Now, to those in Ireland. The best part of doing a PhD here were the people doing their PhDs alongside mine – a wonderfully supportive network of brilliant people from all over the world. I appreciate each and every PhD researcher in the Hardiman Research Building (HRB), Data Science Institute (DSI), Postgraduate Research Society (PGRS), and my PhD program (d-real) for their contribution to creating such a sense of PhD community. In particular, I would like to thank Wei, Aisling, Ed, Yao, Greta, Oksana, Yuchen, Heike, Dhairya, Rimjhim, Andreas, Menna, Maeve M., Charlie, Andrea, David, Ali H., Sina, Nikki, Ashley, Emily, Muz, Pat, Ihab, Dave, Filippo, Victoria, and Divya. Special thanks and love to: thesis writing buddies Ashy, Maria, Maeve F., Heidi, Anne, Rachel, Yaheli, and the PhD pomodoro writing group for their writing solidarity; fellow American and name-sister Sarah for our frequent walk breaks along the River Corrib; Annanda and Felipe for Board Game Sundays (TBC in the Netherlands); HRB buddies Máiréad and Ananya, for their friendship and especially their moral support in the final months of thesis writing; and my dearest friend and kindred spirit Maryam B., for her steadfast friendship and beautiful heart. I would also like to thank Maryam M., Cécile, Maraim, Lukasz (Luke), Augustin, Niall, Dorus, and Jared who, while not (currently) completing PhDs, supported mine with their friendship, lunchtime chats, and one trip to A&E/the ER (thanks, Jared!). I am grateful to call you all my friends, and the hardest part of leaving Galway will be saying goodbye to all of you. Thank you.

Also in Galway, I would like to thank my neighbor (and landlord) Diarmaid, his wife Danuta, and their family, for their kindness and generosity during my time in Ireland. During a cost-of-living and housing crisis, they provided me with a place to live with reasonable rent and treated me with compassion and hospitality. Without them and without finding a place to live, I would not have been able to complete my PhD. I also could not have done this PhD journey without the support of my therapist, yoga teacher, and massage therapist, who helped me overcome the stress and trials of a PhD journey. And, lastly, a thank-you to the city of Galway – to my favorite writing spots, the Full Duck Cafe in Renmore and the Secret Garden in the West End, and to all the birds and wildlife of Ballyloughane Beach for providing me with a refuge to decompress and connect with nature.

And, of course, loads of thanks to my supervisors – Mathieu, Heike, Kathryn, and Dave. I am grateful for your mentorship, guidance, and the opportunity to do this research and grow both professionally and personally. Also thank you to my Graduate Research Committee (GRC) members, John B., Karen, and especially John D., who recommended Killmister’s work to me. Thanks as well to Claire and Ayushi for their feedback on study statistics and design; fellow d-realer Leona for introducing me to Schwartz’s Theory of Basic Human Values; Hardy in the library for Open Scholarship support; Ian in the Technology Transfer Office for helping get the paperwork in order to go build the VcPA in Amsterdam; and to Kevin for his guidance on qualitative research and lending me his books. A special note of gratitude to the philosophy department and the incredible staff at DSI – especially Hilda, Claire, Michelle, and Christiane – for their support navigating University policies, paperwork, and making sure I can travel and conduct the kind of research needed to complete this PhD.

I would like to acknowledge my funder, Science Foundation Ireland’s Center for Research Training in Digitally-Enhanced Reality (d-real), for their financial contribution to this PhD. I am grateful for the generous travel allowance that allowed me to attend conferences and network with colleagues all over the world, and to program manager Stephen, for organizing social events and trying his best to support us during a global pandemic.

Acknowledging this, I must also acknowledge the people who have worked to reform the structural and financial barriers to doing and completing a PhD in Ireland. I would like to thank the Postgraduate Worker's Organization (PWO), formerly PGWA and PCAU. While everyone involved in PWO or its precursors deserves thanks, I would like to explicitly thank a handful of people whose work personally affected me. Locally, I would like to acknowledge Seb, Lennita, and Shane for their tireless advocacy to better the situation of PhDs here in Galway. At the national level, I would like to thank Jeff for his commitment to crafting a better world; Matt for the countless hours he has spent gathering data to advocate for PhD researchers; and Shaakya for standing up for PhD researchers from non-EU countries and marginalized communities. Some of you I have never met in person, but your efforts are inspirational and continue to move me. Thank you, Seb, Lennita, Shane, Jeff, Matt, Shaakya, and the thousands of people fighting for worker recognition, parental leave, holiday leave, a living wage, and fair immigration costs and procedures for PhD researchers.¹

On this note, I would like to again thank my supervisors – in particular, Heike and Mathieu. They have always done their best to support me finishing the PhD, in more ways than one. Not only am I grateful for their academic expertise and professional guidance, but for their empathy, support, validation, and understanding in all aspects of my life. In particular, I would like to thank them for allowing me to travel back-and-forth to my long-distance partner in Germany; for driving me to get a COVID test (thanks, Heike!); for introducing me to Art Nouveau, the École de Nancy, and Émile Gallé (thanks, Mathieu!); and for always creating a space to talk out whatever hurdles – be they structural, bureaucratic, health-related, professional, or otherwise – that I was currently facing.

Internationally, I would like to thank all the people who participated in the research studies that made this PhD possible – in particular, those who participated in the semi-structured interviews, for their honesty, vulnerability, and time. I would also like to thank all the reviewers who provided feedback on my work, especially those at *Digital Society*, and the organizers of Soapbox Science for the opportunity to share my work with the public. Special thanks to thesis proofreaders Seb, Máiréad, Marc, Sarah, Mom, and Dad for volunteering their time. To all the incredible rockers whose songs populate my Apple Music library – thank you for forming the soundtrack of my life and this thesis, and for helping me power through some tough PhD moments. To the examiners, Daniel and Ed, and the chair, Felix, of my viva, thank you for your constructive feedback and for such an intellectually stimulating conversation.

And lastly, I would like to thank the most important person in my life - my love, my best friend, and my constant interlocutor, fellow PhD candidate Marc Barroso Mancha. We did it, love. *Eres mi vida, y no puedo esperar pasar toda una vida contigo.*

*Though I know I'll never lose affection
For people and things that went before
I know I'll often stop and think about them
In my life, I love you more*

The Beatles (“In My Life”)

¹ For more on PWO and these issues, please see: <https://www.pwo-ireland.com>

Abstract

There are a host of data privacy decisions we must make every day – and it is exceedingly difficult, if not impossible, for us to make meaningful decisions about all of them. In this thesis, I define, conceptualize, interrogate, and design for value-centered privacy decision-making – that is, decisions that are focused on who we are and what we value – as a means of respecting and promoting user autonomy. To achieve this, this work utilizes philosophical theory to understand value-centered privacy decisions and translates this theory into a system that promotes such decisions. In summary, this work has **two major contributions**.

Firstly, I **conceptualize and define value-centered privacy decision-making** using a value-centered theory of autonomy. I explore how we can *create the space* for value-centered privacy decisions by applying the Four-Dimensional Theory of Self-Governance (4DT). I first conceptualize privacy decisions in terms of these four dimensions – *self-definition*, *self-realization*, *self-unification*, and *self-constitution* – and explore existing data privacy challenges through this lens. In particular, I conceptualize notice fatigue in terms *self-realization*, *self-unification*, and *self-constitution*; a lack of relevant privacy controls in terms of *self-realization* and *self-unification*; and nudges in terms of *self-realization* and *self-unification*. I then present and discuss results from a mixed-methods investigation into how values are involved in privacy decisions – in particular, app choice. We found that they were related in a highly individualized, context-specific manner, observing different values that were more relevant based on the app in question. This suggests that the value-privacy relationship is largely informed by individual preferences and understandings of values. However, the values of *Use*, *Control*, and *Community* were quite prevalent, with *Use* and *Control* in particular spanning contexts and individual participants. They were also frequently perceived as in conflict with each other. This suggests that these three values are the most relevant to consider when designing for value-centered privacy decisions. The participants’ experiences can also be explained using 4DT, providing empirical support for our conceptualization of value-centered privacy. However, the study results also provide insights into how existing systems – such as surveillance capitalism and the attention economy – frustrate value-centered privacy decisions.

Secondly, I use the 4DT-based understanding of value-centered privacy decisions **to establish the usability and effectiveness of the value-centered approach, designing a privacy assistant** to help users make app choices that are in more accordance with their personal values. To inform the design of a smartphone assistant that creates this space for users, I examine an existing technology – personalized privacy assistants (PPAs) – using the 4DT lens. Using insights from this examination, I propose a value-centered, smartphone privacy assistant (VcPA) to help users make more value-centered decisions at one privacy decision point: smartphone app choices. This VcPA consists of three features: selective notices, exploratory notices, and a “suggest alternative apps” feature. I then present the results from testing a prototype VcPA system with users, serving as a proof-of-concept that a value-centered privacy assistant, designed using privacy preferences *and* values, could help users when making privacy decisions such as choosing apps. In particular, we found that the VcPA prototype helped users download value-consistent apps, with the “suggest alternatives” feature especially well-received. We also identified places where the VcPA could be improved – for example, profiles could be improved by being made more customizable; VcPA notices could be made easier to understand; and the “suggest alternatives” feature could be more streamlined.

This thesis lays the groundwork for future researchers to design systems that promote value-centered privacy decisions. To guide this future work, I lastly present prospective research avenues to advance the value-centered approach to data privacy decision-making. In particular, I discuss limitations of the studies in this work, including engagement with a wider range of demographic groups; touch upon how the identified VcPA improvements, such as improved VcPA profiles, might be accomplished; briefly explore the possibility of applying the value-centered understanding to other privacy contexts; and consider how both system-wide regulation and individual autonomy-enhancing interventions, such as the VcPA, can empower us to shape a technological future based on our values.

Research Outputs Related to this Thesis

Publications

Published

- Doan, X., Florea, M., & Carter, S. E. (2023). Legal-Ethical challenges and technological solutions to e-health data consent in the EU. In P. Lukowicz, S. Mayer, J. Koch, J. Shawe-Taylor, & I. Tiddi (Eds.), *HHAI 2023: Augmenting Human Intellect* (pp. 243–253). IOS Press. <https://doi.org/10.3233/FAIA230088>
- Carter, S. E., Tiddi, I., & Spagnuolo, D. (2022). A “Mock App Store” interface for virtual privacy assistants. In S. Schlobach, M. Pérez-Ortiz, & M. Tielman (Eds.), *HHAI2022: Augmenting Human Intellect* (Vol. 354). IOS Press. <https://doi.org/10.3233/FAIA220212>
- Carter, S. E. (2022). A value-centered exploration of data privacy and personalized privacy assistants. *Digital Society*, 1(27), 1–24. <https://doi.org/10.1007/s44206-022-00028-w>
- Carter, S. E. (2021). Is downloading this app consistent with my values?: Conceptualizing a value-centered privacy assistant. In D. Dennehy, A. Griva, N. Pouloudi, Y. Dwivedi, I. Pappas, & M. Mäntymäki (Eds.), *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics): Vol. 12896 LNCS* (pp. 285–291). Springer International Publishing. https://link.springer.com/chapter/10.1007/978-3-030-85447-8_25

Accepted for Publication

- Carter, Sarah E., & Felzmann, Heike. (2023). How do we value data privacy? Insights and design implications. To be published in: *Engineering and Value Change* (part of: *Springer Philosophy of Engineering and Technology series*). Abstract available at: <https://zenodo.org/record/8367542>

Manuscript in Progress

- Carter, S.E., d’Aquin, M., Spagnuolo, D., Tiddi, I., Felzmann, H., Cormican K. (2023). The privacy-value-app relationship and the value-centered privacy assistant. ArXiv (target journal: *Journal of Business Ethics* (JBE)). <https://arxiv.org/abs/2308.05700>

Conference Presentations and Contributions

- Doan, X., Florea, M., & Carter, S. E. (2023, June 30). Legal-ethical challenges and technological solutions to e-health data consent in the EU. Hybrid Human Intelligence 2023: Augmenting Human Intellect (HHAI2022), Munich, Germany.
- Carter, Sarah E., & Felzmann, Heike. (2023, April 21). How do we value data privacy? Initial results from semi-structured interviews. Forum on Philosophy, Engineering, and Technology (fPET2023), Delft, the Netherlands. Zenodo. <https://doi.org/10.5281/zenodo.8204406>
- Carter, Sarah E., Tiddi, Ilaria, & Spagnuolo, Dayana. (2022, June 13). A "Mock App Store" interface for virtual privacy assistants. Hybrid Human Intelligence 2022: Augmenting Human Intellect (HHAI2022), Amsterdam, the Netherlands. Zenodo. <https://doi.org/10.5281/zenodo.8204393>
- Carter, Sarah E. (2021, September 1). Is downloading this app consistent with my values? Conceptualizing a value-centered privacy assistant. The 20th IFIP Conference e-Business, e-Services, and e-Society (I3E2021). Online. Zenodo. <https://doi.org/10.5281/zenodo.8205147>
- Carter, Sarah E. (2021, July 5). A value-centered exploration of data privacy and personalized privacy assistants. CEPE/IACAP Joint Conference 2021: The Philosophy and Ethics of Artificial Intelligence (CEPE/IACAP 2021). Online. Zenodo. <https://doi.org/10.5281/zenodo.8205315>
- Carter, Sarah E. (2021, June 12). Allowing disclosure: User values and choice in COVID-19 contact tracing applications. Postgraduate Bioethics Conference (PGBC), Institute of Medical Ethics (IME). Online. <https://www.youtube.com/watch?v=edNwZJ1EeM>
- Carter, Sarah E. (2021, February 25). Improving notice: The argument for a flexible, multi-value approach to privacy notice design. 30th Annual Association for Practical and Applied Ethics Conference (APPE 2021). Online. Zenodo. <https://doi.org/10.5281/zenodo.8205297>
- Carter, Sarah E. (2020, December 15). Four-Dimensional autonomy in a digital age: Where are privacy notices going wrong? Ends of Autonomy: December Colloquium. Online. Zenodo. <https://doi.org/10.5281/zenodo.8204502>
- Gleifer, Vaz Alves, Dennis, Louise, Fisher, Michael, Behan, Anthony, Babushkina, Dina, Merdes, Christoph, Archer, Ken, Ní Fhaoláin, Labhaoise, Hines, Andrew, Michael, Loizos, Cardoso, C. Rafael, Ene, Daniel, Evans, Tom, Dennis, Louise, Kaur, Satwant, Carter, Sarah, Grancagnolo, Sergio, & Greidinger, Steven. (2020, June 30). Second Workshop on Implementing Machine Ethics. Online. Zenodo. <https://doi.org/10.5281/zenodo.3938851>

Doctoral Consortium Participation

Fairness, Accountability, and Transparency Conference (FAccT 2022), Seoul, South Korea

The 20th IFIP Conference e-Business, e-Services, and e-Society (I3E2021), Galway, Ireland

Community Outreach and Other Activities

Carter, Sarah E. (2022, November 3-4) A smart office for whom? Workshop speaker and mentor. Write your dystopia (PhD Workshop). LORIA, Université de Lorraine, Nancy, France. <https://members.loria.fr/KFort/teaching/write-your-dystopia-ethics/>

Carter, Sarah E. (2022, October 16). Privacy is(n't) dead: Reclaiming our data in today's online world. Discussion leader. Mount Holyoke College European Alumnae Symposium: Building a More Just Society. The Hague, The Netherlands.

Carter, Sarah E. (2021, September 4). Does this app match my values? Speaker. Soapbox Science Ireland. Galway, Ireland.

Carter, Sarah E. (2021, June 10). Video submission for the Union of Students in Ireland (USI) #WhyResearchMatters Campaign. Online. <https://www.facebook.com/1711669695/videos/10208839184852878/>

Carter, Sarah E. (2021, June 3). Blurring boundaries: Ethics and artificial intelligence. *The Biomedical Scientist*. <https://www.thebiomedicalscientist.net>

List of Abbreviations

4DT	Four-Dimensional Theory of Self-Governance
AI	Artificial Intelligence
DMA	Digital Markets Act
EU	European Union
GDPR	General Data Protection Regulation
HCI	Human-Computer Interaction
IoT	Internet of Things
I-PLOC	Internal Perceived Locus of Causality
MAS	Mock App Store
PPA	Personalized Privacy Assistant
SDT	Self-Determination Theory
SFRF	Selective Facilitated Reflection Framework
SSVS	Short Schwartz Value Survey
SVS	Schwartz Value Survey
TA	Thematic Analysis
TBHV	Theory of Basic Human Values
US(A)	United States of America
UX	User Experience
VcPA	Value-Centered Privacy Assistant
VSD	Value Sensitive Design
WEIRD	Western, Educated, Industrialized, Rich, and Democratic

List of Tables

Table 2-1: The dimensions of 4DT	18
Table 3-1: Understanding data privacy challenges through the lens of 4DT	34
Table 3-2: PPA evaluation using 4DT and suggested modifications for a VcPA	40
Table 4-1: Mixed-methods empirical study design, in three phases.....	48
Table 4-2: Apple Privacy Label options	52
Table 5-1: Features analyzed in survey data.....	70
Table 5-2: Summary of value themes and sub-values from semi-structured interviews.....	79
Table 5-3: Summary of value tensions, with three tensions of particular interest in bold	103
Table 6-1: VcPA feature reception (Likert scale, 1-5).....	125
Table 7-1: Summary of major thesis contributions, findings, and implications for future research.....	131
Table 7-2: Tentative framework, the Selective Facilitated Reflection Framework (SFRF), for VcPA design and deployment in other privacy settings.....	140

List of Figures

Figure 2-1: Self-governance in action according to 4DT	19
Figure 2-2: Schwartz values in a quasi-circular arrangement. Closely related values are in the same color. Modified from: Schwartz (1992), with value definitions from Lindeman and Verkasalo (2010).	23
Figure 3-1: Four-Dimensional Theory of Self-Governance (4DT) as it pertains to individual data privacy decision-making	27
Figure 4-1: Mock App Store main page with a few example apps.....	54
Figure 4-2: Plot of clusters according to <i>Power</i> (x-axis), <i>Achievement</i> (y-axis), and <i>Hedonism</i> (z-axis), where cluster 1 is red, 2 is green, and 3 is blue.....	56
Figure 4-3: Presentation of VcPA profiles to participants on the Mock App Store.....	56
Figure 4-4: Selective notice example	57
Figure 4-5: Exploratory notice example	57
Figure 5-1: Heatmap of significant Spearman correlations ($p < 0.5$) for: total dataset <i>Value</i> , <i>Value App</i> , <i>App</i> , and privacy preference.....	71
Figure 5-2: Differences in <i>Value App</i> ($p < 0.05$, unequal t-test) by t-statistic, where a positive value indicates higher importance to Lose It! and negative to OpenLitterMap	72
Figure 6-1: Number of participants who downloaded profile-matching apps x - y% (in decimal) of the time.....	121
Figure 6-2: VcPA profile reception.....	122
Figure 7-1: Hypothetical example of a selective notice with clearer privacy preference-value mapping and emphasis on values <i>Use</i> , <i>Control</i> , and <i>Community</i>	134
Figure 7-2: Possible future profile design research, with emphasis on profile customizability and the values <i>Control</i> , <i>Use</i> , and <i>Community</i>	135

Chapter 1 Thesis Summary and Structure

*They took the credit for your second symphony
Rewritten by machine on new technology
And now I understand the problems you could see
[...]
Video killed the radio star
Video killed the radio star
Pictures came and broke your heart
[...]
Video killed the radio star
In my mind and in my car
We can't rewind, we've gone too far*

Buggles (“Video Killed the Radio Star”)

Section 1.1 Chapter Overview

In this chapter, I outline the research motivation, aims, contributions of this work, and the structure of this thesis. In summary, this work aims to promote more meaningful, value-centered privacy decisions in a current privacy environment that makes such decisions difficult. Firstly, I use a value-centered theory of autonomy from the philosophical literature to conceptualize and define what it means to make value-centered privacy decisions. I then utilize this understanding to design a prototype smartphone privacy assistant that promotes more value-centered privacy choices. In particular, I look at the decision whether to download a smartphone app. We also conducted a mixed-methods study to assess both the proposed understanding of value-centered privacy decisions and the assistant itself, laying the foundation for future research to build privacy assistants centered on users’ personal values.

Section 1.2 Thesis Summary: Research Motivation, Aims, and Contributions

The current privacy landscape is rich with “dark” design patterns, a never-ending flow of pop-up notices, and impossible-to-read terms and conditions that make managing our own data privacy difficult (Solove, 2021). While much research and discussion has been devoted to the merits of the privacy self-management model itself, I wish to explore the values behind the privacy choices we make. In particular, I wish to identify what values are relevant, how they are involved in our privacy decisions, why we do not act in accordance with them, and how we could promote more value-centered choices. I further put forth that we *should* promote these value-centered choices in the name of respecting user autonomy.

This work is an exercise in bridge-building – it aims to translate philosophical theory into a form that can be implemented by computer scientists. To build this bridge, I define, conceptualize, interrogate, and design for value-centered privacy decision-making. This results in **two major contributions** – one conceptual, and one applied.

Firstly, this work **conceptualizes and defines value-centered privacy decision-making** using a value-centered theory of autonomy from the philosophical literature. I

select and utilize the Four-Dimensional Theory of Self-Governance (4DT) to further define, conceptualize, and understand the relationship between personal values and privacy decisions (Killmister, 2017). I also use this lens to provide insight into why we may not always act in accordance with our values when making data privacy decisions.

Secondly, I use insights derived from 4DT to **design a privacy assistant** to help users make privacy choices that are in more accordance with their personal values. I utilize this 4DT value-centered understanding to inform the design of a proof-of-concept prototype system. This system, called a value-centered privacy assistant (VcPA), builds upon and takes inspiration from a current technology – personalized privacy assistants (PPAs) (Liu et al., 2016). After initially assessing existing PPAs using a 4DT lens, I identify features for a VcPA aimed at promoting value-centered privacy decisions in one context – choosing and downloading smartphone apps. I select this context due to the prevalence of smartphone apps in our lives and literature suggesting *some* relationship between values, smartphones, and/or privacy.

We also conducted a mixed-methods study to evaluate the chosen approach to value-centered decision-making. The study involved an online survey of smartphone users' privacy preferences and values; testing the VcPA prototype; and semi-structured interviews. We use these results to empirically establish how values are involved in data privacy decisions and to assess the 4DT approach. We also use these results to establish the VcPA's effectiveness for promoting value-centered privacy choices and to identify areas for future research.

In summary, this work contributes a novel means of thinking about the role of personal values in privacy, normatively grounded in respect for autonomy. It lays the groundwork for future researchers to build privacy assistants centered on a user's personal values.

Section 1.3 Structure of Thesis

This thesis is organized as follows:

In **Chapter 2**, I present background information pertaining to the value-centered approach to data privacy. I explore the rationale behind and challenges of privacy self-management, before shifting the discussion back to the normative basis of the current self-management model – respecting autonomy. I claim that we should understand respecting autonomy in privacy decisions as promoting value-centered choices – that is, privacy choices that are made in accordance with our personal values. To further define, conceptualize, and investigate this relationship between values and data privacy, I introduce the Four-Dimensional Theory of Self-Governance (4DT) and the Theory of Basic Human Values (TBHV) (Killmister, 2017; Schwartz, 2012).

In **Chapter 3**, I conceptualize value-centered privacy decisions using 4DT (Killmister, 2017). After applying 4DT to privacy decisions to conceptualize value-centered choice, I identify and define three major areas of where we are not making privacy decisions according to our values: notice fatigue, lack of relevant controls, and nudges. I then utilize this understanding to evaluate personalized privacy assistants (PPAs) and inform value-centered privacy assistant (VcPA) design. Based on this analysis of PPAs, I identify three features a VcPA must have to facilitate value-centered privacy decisions: selective notices, exploratory notices, and a “suggest alternative applications” feature.

In **Chapters 4-6**, I present the mixed-methods empirical study. **Chapter 4** presents the study methods, which consisted of three phases. Phase I involved an online survey of values, privacy preferences, and smartphone apps. Phase II involved testing a prototype

VcPA system informed by Phase I results. A testing environment – called the Mock App Store (MAS) – was also designed for testing the VcPA. Phase III consisted of follow-up semi-structured interviews with some Phase II participants. These interviews probed participants’ values, privacy preferences, and app choices on the MAS as well as in their everyday life. The three phases were integrated around two research questions: **RQ1:** *What is the relationship between values and privacy preferences when deciding to download an app, if any?* And: **RQ2:** *How useful and effective is a value-centered privacy assistant at helping users make app choices consistent with their values?*

In **Chapter 5**, I present the results from the online value and privacy preference survey (Phase I) and the relevant interview results (Phase III) to answer **RQ1**. In particular, we explore *how we value privacy*: how values, privacy preferences, and app choices correlate, are understood, and are conceptualized by users. To these ends, the survey provides quantitative insights into the correlations between the 10 general, life-guiding Schwartz values from the TBHV. These quantified results are further contextualized with results from the semi-structured interviews, which provide a deeper, richer understanding of how we value data privacy – including how we describe, present, and understand the role of our values and their tensions. In summary, while values and privacy preferences were related in a highly individualized and context-dependent manner, a few values (*Control, Use, and Community*) were quite prevalent. This suggests that these three values are highly relevant to apps and data privacy decisions. In addition, while 4DT appeared to be a reasonable understanding of value-centered choice based on these studies, the results suggest methodological limitations with using the TBHV to operationalize values.

In **Chapter 6**, I report the results of the Mock App Store Study (Phase II) and follow-up semi-structured interviews to answer **RQ2**. In particular, we use the results to evaluate the desirability and effectiveness of the proof-of-concept VcPA system. In summary, the VcPA prototype helped users download value-consistent apps, with the “suggest alternatives” feature especially well-received by participants. Considering participant feedback, we also identified places where the VcPA could be improved – for example, profiles could be improved by being made more customizable; VcPA notices could be made easier to understand; and the “suggest alternatives” feature could be more streamlined.

In **Chapter 7**, I conclude and present future avenues for the value-centered approach to data privacy decision-making. After an overview of the major research findings, I discuss the limitations of the studies conducted in this work, including the inclusion of a broader range of demographic groups. I also discuss the means by which the identified VcPA enhancements, such as improved VcPA profiles, could be achieved. Additionally, I briefly consider the application of the value-centered understanding to other privacy contexts. Lastly, I consider the role of both system-wide regulation and individual autonomy-enhancing interventions, such as the VcPA, in empowering us to construct a technological future that aligns with our values.

Chapter 2 Background

*All I want is to be left alone, in my average home
But why do I always feel
Like I'm in the Twilight Zone?
I always feel like somebody's watchin' me
And I have no privacy
I always feel like somebody's watchin' me
Tell me, is it just a dream?*

Rockwell (“Somebody’s Watching Me”)

Section 2.1 Chapter Overview

In this chapter, I set the stage for a value-centered approach to data privacy. I begin by exploring the rationale behind and challenges of privacy self-management – cognitive biases, heuristics, nudges, and (bright and dark) design patterns. To address these challenges, I propose a return to the normative foundation of this self-management model – respect for autonomy – based on human values, understanding respect for autonomy as *promoting value-centered decisions*. To more deeply investigate this relationship between values and privacy, I draw upon the Four-Dimensional Theory of Self-Governance (4DT) and operationalize values using the Theory of Basic Human Values (TBHV) (Killmister, 2017; Schwartz, 2012).

2.1.1 Collaborator Contributions

The ideas described in this chapter are my (the PhD candidate’s) own work. Feedback was provided by PhD supervisors Dr. Heike Felzmann, Prof. Dr. Mathieu d’Aquin, Prof. Dr. Kathryn Cormican, and Dr. Dave Lewis.

2.1.2 Relevant Papers and Conference Contributions

Some material in this chapter, including certain text and figures, has been previously published or presented in the following:

Carter, S. E. (2022). A value-centered exploration of data privacy and personalized privacy assistants. *Digital Society*, 1(27), 1–24
<https://doi.org/10.1007/s44206-022-00028-w>

Carter, S. E. (2021). Is downloading this app consistent with my values?: conceptualizing a value-centered privacy assistant. In D. Dennehy, A. Griva, N. Pouloudi, Y. Dwivedi, I. Pappas, & M. Mäntymäki (Eds.), *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics): Vol. 12896 LNCS* (pp. 285–291). Springer International Publishing. https://doi.org/10.1007/978-3-030-85447-8_25

Background

Carter, Sarah E. (2021, September 1). Is downloading this app consistent with my values? Conceptualizing a value-centered privacy assistant. The 20th IFIP Conference e-Business, e-Services, and e-Society (I3E2021). Online. Zenodo. <https://doi.org/10.5281/zenodo.8205147>

Carter, Sarah E. (2021, September 1). PhD proposal: Conceptualizing and realizing a value-centered privacy assistant. Doctoral Symposium: The 20th IFIP Conference e-Business, e-Services, and e-Society (I3E2021), Galway, Ireland. Zenodo. <https://doi.org/10.5281/zenodo.8204916>

Carter, Sarah E. (2021, July 5). A value-centered exploration of data privacy and personalized privacy assistants. CEPE/IACAP Joint Conference 2021: The Philosophy and Ethics of Artificial Intelligence (CEPE/IACAP 2021). Online. Zenodo. <https://doi.org/10.5281/zenodo.8205315>

Carter, Sarah E. (2021, February 25). Improving notice: The argument for a flexible, multi-value approach to privacy notice design. 30th Annual Association for Practical and Applied Ethics Conference (APPE 2021). Online. Zenodo. <https://doi.org/10.5281/zenodo.8205297>

Carter, Sarah E. (2020, December 15). Four-Dimensional autonomy in a digital age: Where are privacy notices going wrong? Ends of Autonomy: December Colloquium. Online. Zenodo. <https://doi.org/10.5281/zenodo.8204502>

Section 2.2 The Breakdown of Privacy Self-Management

Those of us browsing the Web from the European Union have become accustomed to a familiar sight – privacy notices, or “privacy pop-ups,” asking us to agree to cookies. These notices are a result of the ePrivacy Directive² and the General Data Protection Regulation (GDPR), the latter of which was passed in 2016 and implemented in 2018 (General Data Protection Regulation (GDPR), 2016; EPrivacy Directive, 2009). Under the GDPR, informed consent is one of six legal bases³ by which personal data can be collected and processed. This legal basis resulted in a flurry of privacy notices to elicit explicit, unambiguous consent (Degeling et al., 2019).

While disclosures of this kind existed before the GDPR, the sheer volume of them has brought discussions surrounding the value and effectiveness of privacy self-management back to the forefront. Privacy notices and other privacy self-management mechanisms are preceded by a rich history of consent in other fields, especially bioethics.⁴ Like consent forms for medical procedures or clinical trials (Beauchamp, 2011), privacy

² To be replaced by the ePrivacy Regulation, currently under negotiation by the legislative bodies of the European Union (the EU Commission, Parliament, and the Council of Ministers). A statement issued by the European Data Protection Board regarding the new legislation can be viewed here: https://edpb.europa.eu/system/files/2021-03/edpb_statement_032021_eprivacy_regulation_en_0.pdf

³ The six legal bases for processing data are: consent; contract; legal obligation; vital interests; public task; or legitimate interests. These are specified in Article 6 of the GDPR.

⁴ There are special considerations to medical research that may not be applicable to all data collection situations – e.g., issues of dependency and vulnerability in healthcare situations which introduce added consent concerns. There are, however, overlaps between bioethics and data ethics - especially as Big Data, smartphone apps, and artificial intelligence enter the medical field. For examples, see: Jongsma et al. (2018), Klugman et al. (2018), or Lucivero & Jongsma (2018).

Background

notices, terms and conditions, and other privacy self-management disclosures should allow individuals to make reasoned, informed choices concerning their data privacy.

2.2.1 “Uniquely Human”: Why Data Privacy is Not Dead

Asking for consent has normative roots in respecting autonomy. Foundational documents of bioethics, such as the Belmont Report and the Declaration of Helsinki (K. J. Ryan et al., 1979; World Medical Association, 2013), stress the importance of respecting autonomy as a means to checking exploitative and manipulative practices.⁵ Similarly, online or on our smartphones, we face risks of manipulation and surveillance. These risks are intimately related to existing incentives promoted by our economic and governmental structures.

One need only look to the work of Shoshana Zuboff to see these risks. In her work, *The Age of Surveillance Capitalism: The Fight for a Human Future at the Frontier of Power* (2019), Zuboff interrogates the business models of large technology companies (such as Google (Alphabet), Facebook (Meta), and Amazon) and their worrying uses of data. Following the collapse of the .com bubble in the early 2000s, Zuboff suggests that Google founders Larry Page and Sergey Brin discovered that personal data – such as search history and clicks, formerly viewed as nothing more than *data exhaust* – was incredibly valuable in terms of predicting a person’s future behavior (Zuboff, 2019). This created what Zuboff refers to as a *prediction marketplace*, where data is sold to third parties and processed by progressively complex algorithms to predict and shape our futures.

Originally utilized for targeted advertising and predicting online consumer behavior, Zuboff further describes how this data has been increasingly utilized to influence *real-world* behavior. Some of these uses, such as the use of personal data to influence a person’s political actions by Cambridge Analytica,⁶ can be considered outright manipulative.⁷ The Cambridge Analytica scandal was a particularly significant case because of the kind of influence it entailed; it marked a clear break from market-motivated influence (e.g., profit-making) to politically motivated influence (e.g., election influence).

In this vein, it is also important to note that the risks of data collection are not limited to corporate uses and repercussions, but government ones as well.⁸ For example, following 9/11, “The Patriot Act” was passed in the United States, resulting in expanded government surveillance measures in the name of national security (Ombres, 2015). The extent of this surveillance was famously exposed by whistleblower Edward Snowden in 2013. In *Permanent Record* (2019), Snowden details a program of mass surveillance that included the average US citizen’s personal communications – including personal phone calls, text messages, and emails.⁹ When describing this post-9/11 expansion of citizen

⁵ The traditional catalysts for the Belmont Report and the Declaration of Helsinki were the Tuskegee syphilis trials, where black men were left with untreated syphilis and observed, and revelations surrounding Nazi experimentation on prisoners during World War II. See: Brandt (1978) and Carlson et al. (2004).

⁶ See a collection of reporting by *The Guardian* at: <https://www.theguardian.com/news/series/cambridge-analytica-files>

⁷ Here, I use Daniel Susser and colleagues’ (2019) definition of online manipulation as an action that uses a person’s individual decision-making vulnerabilities – gleaned from personal data collected online, and often subconscious – to intentionally and covertly influence a user to make a pre-determined decision (pg. 4). These decision-making vulnerabilities are detailed in Section 2.2

⁸ I will be focusing on democratic governments here given this work’s mostly European (with some US) privacy framing. For those interested in issues surrounding surveillance and authoritarian governments, see: Deibert (2015).

⁹ In addition, governments are heavily dependent on tech companies to boost their own ability to surveil. A close relationship between Google and the US government has previously gained media attention – for example, under the Obama Administration, then-Google head and former Obama campaign adviser Eric

Background

surveillance, privacy experts frequently refer to philosopher Jeremy Bentham's panopticon – a Victorian-era method of designing prisons with a guard centrally located so prisoners are constantly being observed.¹⁰ However, unlike Bentham's panopticon, constant collection and storage of one's data – *dataveillance*, to use the term from Simon (2005) – continues to provide more and more access to one's life as technology progresses. Centralized dataveillance weakens civil liberty protections by creating large datasets that could be accessed by law enforcement (Shackleton, 2019). They also produce a *chilling effect*, or the tendency of individuals to watch what they say or share online when they are aware of being surveilled (Solove, 2006).¹¹

While the chilling effect itself is a concern for democracies, a shrinking private sphere could have repercussions for the development of autonomous, reflective citizens (Cohen, 2013; Zuboff, 2019). Zuboff (2019) draws upon research in developmental adolescence to further suggest that surveillance capitalism is stunting the ability of today's adolescents to mature into adulthood. Instead of learning to balance their public and private sphere and develop a sense of self independent of outside influence (“I think,” “I feel,” “I believe” (Zuboff, 2019, pg. 454)), young people are increasingly subjected to “Life in the Hive” (Zuboff, 2019, Chapter 16) – an addictive network of likes and curated feeds, constantly subjected to social pressure.

“What are the consequences of the failure to win a healthy balance between inner and outer, self and relationships? Clinical studies identify patterns associated with development stagnation. Not surprisingly, these include an inability to tolerate solicitude, the feeling of being merged with others, an unstable sense of self, and even an excessive need to control others as a way of keeping the [social media] mirror close” (Zuboff, 2017, pg. 455).

According to Zuboff, “Life in the Hive” has repercussions for democratic institutions – if one remains in a state of perpetual adolescence and does not possess a fully-formed sense of self, it is challenging to assert one's interests and participate effectively in a democratic system. The loss of this reflective space challenges something *uniquely human* – *our ability to shape our lives according to our values*. “Life in the Hive” or constantly under surveillance is simply not natural for human beings. Just as “industrial capitalism depended upon the exploitation and control of nature,” surveillance capitalism depends upon “the exploitation and control of human nature” though the use of our personal data (Zuboff, 2019, pg. 470).

At the center of the manipulative, human-nature-shaping potential of data, of course, is Artificial Intelligence (AI). AI-backed personalized, dynamic, and seductive targeted advertising is eroding our autonomy, further exacerbating the already questionable tactics to bypass conscious thought used in traditional advertising (Susser et al., 2019; Yeung, 2017). AI-generated social media newsfeeds create “silos” that make it especially

Schmidt was reported to have met some 427 times with President Obama between 2009 and 2015 (Dayen, 2016; Estes, 2011). On this note, Julian Assange, famous founder of WikiLeaks and well-known for his controversial free press activism, writes that “there is a comfortable willingness among privacy campaigners to discriminate against mass surveillance conducted by the state to the exclusion of similar surveillance conducted for profit by large corporations. [...] The movement to abolish privacy is twinned-horned. Privacy advocates who focus exclusively on one of those horns will find themselves gored by the other” (Assange, 2016).

¹⁰ For example, see: Campbell & Carlson (2002) and Simon (2005).

¹¹ Supporting this idea, a 2019 study empirically demonstrated that awareness of online government surveillance resulted in a decrease of online political participation at the height of the 2016 US election season (Stoycheff et al., 2019).

Background

easy to manipulate through tailored, personalized user content. Tristan Harris, founder of the Center for Humane Technology¹² and former Google employee, describes what he calls a “race to the bottom of the brain stem” where companies compete for user attention through personalized data-fed recommender systems to “upset us, polarize us, and addict us” (Harris, 2020). This “attention economy” not only fuels societal polarization, but further utilizes subversive tactics aimed at targeting our reactive, “brain stem” selves over our higher cognitive capacities (Davenport & Beck, 2002; Goldhaber, 1997).

Besides helping create our own personalized addictive silos, AI can also cause harm through AI bias. The harms from AI bias disproportionately affect marginalized communities and intensify old disparities due to the added speed of the algorithms.¹³ For example, data-fueled AI has been used by governments for social benefit or welfare fraud detection. While the most obvious potential harm of using AI in these instances is, of course, inaccurate data resulting in inaccurate data predictions, AI bias is often much more insidious (D’Ignazio & Klein, 2020; O’Neill, 2016). The childcare benefits scandal in the Netherlands¹⁴ is a telling representation of this point. Between 2013 and 2016, an estimated 26,000 parents were wrongfully accused of fraud due to an algorithm utilized to identify fraud for childcare benefits (Henley, 2021). The algorithm rated dual citizens or those who were born outside the Netherlands as higher risk, leading to a disproportionate number of those with immigrant backgrounds being falsely flagged (*Dutch Childcare Benefit Scandal an Urgent Wake-up Call to Ban Racist Algorithms*, 2021; Henley, 2021). Parents were asked to pay back tens of thousands of euros, leading to financial hardship, more than 1000 kids taken from their homes, divorce, and depression (“1,675 Children Removed from Parents’ Custody in Benefits Scandal,” 2022; *Dutch Childcare Benefit Scandal an Urgent Wake-up Call to Ban Racist Algorithms*, 2021; Henley, 2021). Tellingly, an investigative committee described the affair as “een ongekend onrecht” (an unprecedented injustice) (“Commissie: Ongekend Onrecht in Toeslagenaffaire, Beginselen Rechtsstaat Geschonden,” 2020).¹⁵

Taken together, a shrinking private sphere *is dangerous* – it threatens our civil liberties; our democracies; our futures; and our ability to form a sense of self.

2.2.2 The Illusion of Homo Economicus

Privacy self-management and its manifestations, including privacy notices, is one line of defense against these risks. James F. Childress, co-author of the Belmont Report, further proposes that respect for autonomy requires that we not only refrain from manipulating (a negative duty), but also that we disclose information and foster autonomous decisions

¹² See: <https://www.humanetech.com>

¹³ For a plethora of examples of AI bias and discrimination, see the Algorithmic Justice League at: <https://www.ajl.org>

¹⁴ *De toeslagenaffaire* or *de kinderopvangtoeslagaffaire* in Dutch. For reporting on the scandal, see (in Dutch): <https://nos.nl/collectie/13855/artikel/2364513-kabinet-rutte-iii-gevallen-wiebes-helemaal-weg> or (in English): <https://www.theguardian.com/world/2021/jan/14/dutch-government-faces-collapse-over-child-benefits-scandal>

¹⁵ In 2023, the Netherlands was in the news again. A joint investigation by Lighthouse Reports, WIRED, and Vers Beton into the algorithm previously used in Rotterdam to flag potential welfare fraud suggests it was very biased. For example, one of the biggest determinants of being flagged as high risk was Dutch proficiency, leading to, again, immigrants being disproportionately targeted. Rotterdam officials have since abandoned plans to build and deploy a new algorithm. Instances of AI bias such as these are happening worldwide – e.g., footnote 13. While these cases in the Netherlands are disturbing, it is commendable that Rotterdam handed over their model for investigation and public scrutiny, something the investigators themselves note in their report. For more on this story, see: <https://www.lighthousereports.com/investigation/suspicion-machines/>

(positive duties) (Childress, 1990). These duties have been operationalized into *informed consent* – that is, facilitating voluntary choice by providing information that is comprehensive and complete (Beauchamp, 2011). For privacy self-management to be an effective tool for respecting autonomy, then, individual choices should be voluntary and based on comprehensive information.

Discussions surrounding the “privacy paradox” phenomenon – or the observation that the choices we make when faced with data privacy decisions do not match what we say our privacy preferences are – suggests that neither of these criteria for informed consent are met (Spiekermann et al., 2001). There are two general arguments that have been used to explain the privacy paradox, one in line with privacy decision making as a rational, preference-based endeavor and the other postulating that humans are “flawed” thinkers who make questionable privacy decisions (Solove, 2021).

The first claims that our privacy behavior, and not our stated privacy preferences, *truly* reflects how much we value privacy against, say, the good or service being offered. This is because we have been adequately informed by the privacy notice or policy and are acting in a rational manner. This view has roots in American privacy scholar Alan Westin’s segmentation model, where a rational consumer reads privacy policies and makes decisions based on their preferences (Hoofnagle & Urban, 2014). In this model, there are three types of consumers: “privacy pragmatists,” “privacy fundamentalists,” and “privacy unconcerned,” with the pragmatists in the majority.

In contrast, the second camp claims that our data privacy choices are not an accurate reflection of privacy’s value but rather the result of human bias and heuristics – cognitive “tricks” or shortcuts that help us make fast decisions (Thaler & Sunstein, 2008; Tversky & Kahneman, 1974). Hoofnagle & Urban (2014) explore the influence and pitfalls of Westin’s segmentation model, proposing that the majority of “pragmatist” users are hardly pragmatic at all, often making decisions with little or no understanding of the privacy protections in place. Cognitive and behavioral psychology and its offshoot, behavioral economics, overwhelmingly supports this latter view – and, indeed, few would agree that we are purely rational *homo economicus*¹⁶ (to use Hoofnagle & Urban’s term) when it comes to the decisions we make when faced with a privacy notice.¹⁷ Not only, then, are we not informed enough to manage our data privacy, but we are also susceptible to manipulation through exploitation of heuristic-based thinking. To explore these challenges in more depth, four examples of this “flawed” human thinking, followed by a brief look at the lack of comprehension problem, are explored below.¹⁸

To begin, there are *framing effects*. We are more likely to consent to data collection if it is framed in nonthreatening ways (O’Neill, 2002; Thaler & Sunstein, 2008). In addition, we can be encouraged to consent by utilizing the *inertia bias* – or, the human tendency to stay with the default condition (Thaler & Sunstein, 2008). For example, pre-selecting “Agree to All” in privacy notices increases consent to online data tracking (Utz et al., 2019).

We can also be coaxed into a false sense of security by utilizing the *representative heuristic* and/or the *conjunction fallacy* (Lewis, 2017; Thaler & Sunstein, 2008; Tversky & Kahneman, 1974). Privacy policies are usually long, complicated documents packed with legal terminology (Jensen & Potts, 2004). We may believe something to be true because it matches our mental images of what it should look like – in this case, a long legal document resembles a “strict” privacy policy. This is the *representative heuristic*. Closely related to

¹⁶While the Westin model has its issues, I will return to it briefly in Chapter 3 as an initial conceptual tool.

¹⁷ There is also discussion on whether privacy is something that can or *should* be tradable for a good or service. For example, see: Allen (2013).

¹⁸ For a more comprehensive overview, see: Solove (2021).

Background

the representative heuristic is the *conjunction fallacy*, or the tendency to have our predictions misled when flooded with truthful facts (Lewis, 2017; Tversky & Kahneman, 1974). Privacy policies can be written to selectively reveal or flood us with truthful facts in a way that can mislead our predictions because they remain tactfully vague on other aspects.

Indeed, service providers that collect data can utilize these design strategies to coax us into making the choices they want us to make. Even the simplest privacy notices can be designed to push users towards greater information disclosure, by, say, highlighting the “Accept All” button.¹⁹ These designs, called “dark patterns” or “deceptive design patterns,” are ubiquitously used on privacy notices and have noticeable, measurable effects on data privacy decision-making (Brignull, n.d.; Gray et al., 2018; Mathur et al., 2019; Utz et al., 2019). To describe the consent given on these privacy notices as voluntary would be a stretch. While the GDPR sets out guidelines for what can be considered valid consent – that is, that it must “be freely given, specific, informed, and unambiguous” [Rec.32, Art.2-11], it was estimated that nearly 90% of pop-up notifications do not meet these minimum requirements (Nouwens et al., 2020).

Regarding comprehension, while it is difficult to measure and results vary, it has been suggested that as few as 0.24% of us read online privacy policies (Jensen & Potts, 2004). Practically speaking, no reasonable person *could read all of them* – one study estimated that it would take 244 hours to read all the privacy policies we see in one year alone (McDonald & Cranor, 2008). Comprehension may be difficult due to this informational overload or complex language (Derguech et al., 2018). There is also the well-know “transparency paradox,” where providing too much information is overwhelming while too little could result in something especially important being left out (Nissenbaum, 2011).²⁰ This is further complicated by AI, where much of our data is eventually fed to. How to best make AI transparent, understandable, and explainable remains an area of active cross-disciplinary exploration and debate.²¹

Like the challenge of *privacy policy fatigue*, too many privacy notices can cause *privacy notice fatigue*. This results in many of us deciding to simply “click through” notices rather than reading them (Ben-Shahar & Schneider, 2010; Schaub et al., 2015). Disturbingly, a study of online social networking services estimated that this “click through” rate could be as high as 74% (Obar & Oeldorf-Hirsch, 2018). Eventually, this notice fatigue can cause us to become (what I am calling) “apathetic users” – those who decide to consent every time to a service’s data collection practices because they “no longer care” about their data privacy. The term “apathetic user” used here is meant to capture those who would prefer to be more data protective but feel overwhelmed by notices to the point of “no longer caring,” or apathy. The purpose of this term is to capture a state in which an individual 1.) has internalized apathy and 2.) takes value-inconsistent action. This contrasts with someone who *genuinely* does not care about privacy, where “clicking through” privacy notices would be a value-consistent action.²²

¹⁹ There is a wonderfully informative “Wall of Shame” of such dark or deceptive patterns, encompassing more than just privacy notices (e.g., unsubscribing from an online service). See: <https://www.deceptive.design/hall-of-shame>

²⁰ Privacy notice-and-consent requires providing complete information, but it is often overwhelming. Attempting to abbreviate all this information, however, likely means the loss of critical details, such as: “who are the business associates and what information is being shared with them; what are their commitments; what steps are taken to anonymize information; and how will that information be processed and used” (Nissenbaum, 2011, pg. 36).

²¹ For an overview of AI transparency and explainability challenges from cross-disciplinary perspectives, see: Ferrer et al. (2021).

²² This nuance will be explored in greater detail in Chapter 3.

Background

We are, in summary, overwhelmed and coaxed into making choices that may not be in our best interest instead of informed by privacy notices, privacy policies, and other privacy self-management disclosures.

2.2.3 To Nudge or Not to Nudge: That is the Question

These revelations from behavioral and cognitive psychology pose significant threats to the value of privacy notices and privacy self-management more broadly. To address these challenges, research has been conducted to redesign and improve privacy notices to be more effective. Nearly every aspect of notice design – from timing, channel, modality, and control – has been extensively explored (Schaub et al., 2015). As privacy scholar Ari Ezra Waldman writes, privacy notices and policy designers are “like painters who use line, color, contrast, and perspective to help guide their audiences through a visual narrative” (Waldman, 2018, pg. 127).

Particular interest has been paid to utilizing our heuristics to our benefit – that is, to “nudge” us beneficially. Nudging involves altering someone’s environment in such a way that encourages one decision, without forbidding alternative decisions (Thaler & Sunstein, 2008). In the case of privacy notices, these nudges can be, for example, in the direction of better privacy (preserving) choices, or to increase comprehension/retention of a privacy policy (Almuhimedi et al., 2015; Calo, 2012; Kelley et al., 2013; Waldman, 2018). These interventions have ranged from presenting privacy permissions in a clearer format (“Privacy Facts”) when downloading a new smartphone application (Kelley et al., 2013) to informing users how much their location has been used (Almuhimedi et al., 2015). These privacy-preserving nudges on privacy notices are sometimes called, to contrast them to dark patterns, “bright patterns” (Grassl et al., 2021).

By utilizing notices to encourage us to “preserve our privacy,” however, bright patterns can undermine the original goal of privacy notices – to respect autonomy and prevent manipulation. This is due to the nature of nudges themselves. There remains a fundamental, ongoing discussion concerning whether nudges are inherently paternalistic or manipulative, linked to larger discussions about autonomy. Of particular interest is the concept of volitional autonomy – or that our actions should reflect our authentic desires. Volitional autonomy can be traced to analytical philosophers Harry Frankfurt and Gerald Dworkin. According to Frankfurt’s original theory of free action, to act freely – autonomously – we must act in accordance with our second-order volition, or, in other words, we must desire what we desire (Frankfurt, 1971). Dworkin similarly stated that autonomy is our ability to reflect and endorse “first-order” desires in accordance with our “higher-order” values (Dworkin, 1988b). The “nuder” – the person who is implementing the nudge – is encouraging the actions they desire. They are not necessarily what the person being nudged desires, and the nudged person may act in a manner that they do not endorse (Hausman & Welch, 2010; Schmidt & Engelen, 2020). Put a bit more clearly, bright patterns encourage us to preserve our privacy even in instances where we do not “truly” want to. This, proponents argue, violates our autonomy.

This concern around nudges and user autonomy more generally has been raised in other fields. In the field of psychology, for example, a similar conception of autonomy to that of Frankfurt’s and Dworkin’s has been operationalized by psychologists Richard Ryan and Edward Deci into their Self-Determination Theory (SDT). SDT defines autonomy as self-endorsement of an action according to one’s values (R. M. Ryan & Deci, 2004). Empirically, studies utilizing SDT’s Internal Perceived Locus of Causality, or I-PLOC, measurement suggest that even small nudges can undermine someone’s autonomy (Arvanitis et al., 2020). This has implications for bright patterns, whose use – however

well-intentioned – could cause an I-PLOC shift, again violating, rather than respecting, autonomy.

We are stuck at an impasse when it comes to how to design for privacy self-management. Insights from behavioral psychology suggest that informed consent via notice is not plausible. Instead, we are at worst, uninformed, and at best, overwhelmed when encountering a notice in our current privacy environment. We are manipulated and framed to “consent” our data away. Some have aimed to use privacy notices to nudge us to make “better” privacy choices, but this has possible negative implications on our individual autonomy. Perhaps, as many scholars argue, we should not use privacy notices at all.

Section 2.3 Personal Values: From Self-Management to Self-Governance

I would instead argue that we could try to imagine a role for privacy notices outside that of eliciting informed consent. Perhaps it is not the notices themselves that are the problem, but a mismatch between the task and the tool. This argument has been made by Daniel Susser (2019), who suggests that privacy notices could perhaps promote individual decision-making by increasing awareness of the current data privacy landscape, helping identify potential privacy concerns, and assessing our legal rights, where applicable.²³ I would like to take this one step further, exploring notice outside traditional notice-and-consent by returning to notice’s original goal of respecting autonomy.

To respect autonomy, then, we could try to utilize the insights from psychology and accounts of autonomy in philosophical literature to understand autonomy in a manner that does not rely upon informed consent. I am particularly interested in conceptualizing autonomy in a manner that better captures the role of personal values in data privacy decisions—what I call value-centered privacy decisions. Such an autonomy-focused approach may still be liable to the critiques of privacy self-management approaches – such as the limit of individual approaches to data privacy in the face of surveillance capitalism and other system-wide structures. However, respecting autonomy is still a worthwhile aim for data privacy decisions. Respecting autonomy in data privacy decisions still provides a critical check to exploitation and manipulation despite informational power asymmetries – it is, at the very least, the first line of defense (Susser et al., 2019). In addition, there is also a sense of disempowerment and an increasing learned norm of simply “giving up” on protecting our data privacy. In the case of those of us who reside in liberal democracies, this disempowerment is fundamentally out-of-synch with the central tenets of our governments.²⁴ While we must not prioritize respect for autonomy above all else and still balance it with other values relevant to data privacy – values such as accountability, transparency, and trust (O’Neill, 2020; Waldman, 2015, 2018),²⁵ the hope is such an approach focused on individual, personal values could serve as a complement to these approaches that consider broader norms and values.

²³ See Susser (2019, pg. 38): “If the problem with notice-and-consent as a whole is that it fails to facilitate and respect individual agency over data, then we ought not to deprive ourselves of even flawed and partial mechanisms for strengthening such agency.”

²⁴ I particularly appreciate how this is stated in Susser et al., 2019, pg. 8: “Autonomy is in many ways the guiding normative principle of liberal democratic societies. It is because we think individuals can and should govern themselves that we value our capacity to collectively and democratically self-govern.”

²⁵ Trust—and placing it well—is a pressing concern for data privacy decisions. For example, Onora O’Neill (2020) outlines several problems with trust and accountability in the digital age. Technology allows new intermediaries to control online content opaquely, and it is difficult to decide whether they are trustworthy. This results in unrealistic accountability mechanisms, like data privacy policies and excessive notices.

2.3.1 Autonomy-Values-Privacy: The “Digital Home” Analogy

I see value-centered privacy decisions as those that result in ends that accurately reflect our personal values - a relationship between autonomy, values, and privacy. This idea can be captured in terms of a “digital home.” The furniture you have there; the art you put on the wall; and who you let in should ideally reflect who you are and what you value. We can imagine, however, instances when this would not be the case – perhaps your local furniture store does not carry the color of furniture you like and other stores are too far away for you to reasonably travel to. Similarly, your smartphone or computer is like your “digital home” – which apps you download and who you allow access to your data should ideally reflect your values.²⁶ However, hurdles such as dark patterns may nudge you towards privacy choices do not reflect your values. Individual privacy decisions can therefore be evaluated to the extent to which they result in ends that accurately reflect the individual’s personal values. This intuition is also backed by multidisciplinary literature, which I will survey in the following sections.

a Focus: Personal Autonomy and Personal Values

First, though, I need to present a few caveats – what value-centered privacy decisions are meant to encompass, and what they are not.

Notably, because of the emphasis here on *personal* values, I will not be surveying nor addressing the implications of privacy disclosure or the social value of privacy.²⁷ Similarly, value-centered privacy decisions are not meant to encapsulate discussions about broader, public values that may be relevant for governments when constructing policies with data privacy implications.²⁸ I am not aiming to make regulatory recommendations or assess the GDPR, which, as this chapter demonstrates, likely will need to atone to the lessons of psychology as well. I leave this task to other scholars.²⁹

I am also primarily interested in accounts of *personal autonomy*. This is because I am not aiming here to assess whether our decision to disclose our data affects our overall autonomy or greater democratic processes, but rather, how autonomous the decision to disclose our data is.³⁰ I am not investigating whether what someone values and acts on in their privacy decision is the most morally justifiable one (*moral autonomy*), but whether it is a reflection of their values.

I do not intend to dismiss these (important!) discussions, but rather to complement them with one re-focusing on the individual experience of data privacy choices—those

²⁶ As explored more deeply in upcoming section (b), privacy “the value” is here understood as an instrumental value, valuable for other terminal values that making a privacy decision (i.e., being private or not) brings about. This allows us to consider not only traditional values associated with privacy-preserving behavior (such as *security*), but also other values that may be associated with sharing data (such as the *connection* value of sharing photos on Facebook to keep family and friends “up to date” on your life).

²⁷ For an example of a discussion exploring the social value of privacy, see: Roessler & Mokrosinska (2013).

²⁸ Examples of these discussions can be found around COVID-19 tracking applications. The proposed benefit to public health was considered in terms of: efficacy and uptake of the apps (Luciano, 2020; Morley et al., 2020); justice, equity, and solidarity with vulnerable subpopulations (Hendl et al., 2020); and civil liberties and surveillance (Kitchin, 2020), to name a few.

²⁹ This has been done quite well elsewhere—in particular, Frederik Zuiderveen Borgesius devoted an entire PhD to this subject. While this was completed in pre-GDPR (EU Data Protection Directive) days, many of his analyses regarding informed consent remain relevant to the GDPR. See Zuiderveen Borgesius (2015).

³⁰ For examples of these discussions see: Cohen (2013) who discusses the role of privacy in setting boundaries from external influences, allowing us to define who we are and what we want; or Zuboff (2019) who explores the dangers disclosing data and behavioral modeling pose to our agency and the greater democratic process.

small, often unsatisfying privacy decisions that we make daily as we interact with digital technologies.

b Privacy and Values

Many disciplines have explored the relationship between privacy and values more broadly, albeit in different contexts and with different emphases, and we can see aspects of this captured in the “digital home” analogy.

On the legal side, Daniel Solove has conceptualized privacy (more generally) as many related items that can be encompassed under a common heading without necessarily having a single theoretical basis (Solove, 2002, 2007). He further puts forth that privacy is instrumental in that it allows an agent to protect or promote valuable ends (Solove, 2002). Seeing individual privacy decisions as value reflection retains his idea of both value plurality and instrumentality. Here, however, value reflection emphasizes less the ends that are brought about than which personal values are ultimately reflected by the user achieving those ends.

Values as ends-in-themselves is drawn partly from Value Sensitive Design (VSD), which proposes that technology embeds and expresses values (Friedman et al., 2008). I say in part because we are understanding value-centered privacy decisions centered on a user’s personally held values as something that can be *designed for*. We recognize that technology is not neutral and that design choices are intimately related with human values. In this case, our focus is on one’s personal values rather than greater stakeholder values that concern a designer utilizing the classic VSD approach; it is, one could say, VSD with an emphasis on the individual experience. While we are not engaging in the traditional tripartite methodology associated with VSD (Friedman et al., 2008), we theoretically draw from its understanding of technology as inherently value-laden.

Value-centered privacy decision-making also incorporates an aspect of *information flow*, or, in this case, the data that is being collected as the result of a user’s privacy decision. This is a central idea of philosopher Helen Nissenbaum’s contextual theory of privacy (Nissenbaum, 2004), which puts forth that privacy can be understood as *appropriate information flow* that match context-specific norms. Like VSD, we are zooming in to focus on personally held values over generally consensual norms. Instead of looking for the appropriateness of (data) flows according to norms, we are interested in looking at the appropriateness of flows according to an individual’s specific value set following a data privacy decision. Value-centered privacy decisions encompass those decisions made with due consideration of one’s value set, and the resulting flow of shared data is in alignment with these values.³¹

³¹ Whether the resulting flow of data would be considered appropriate according to generally held norms is not being considered. We can think of instances where a data privacy decision could be value-centered but conflict with greater norms because the individual’s value set does not align with what is considered broadly appropriate. For example, consider a person facing a difficult medical diagnosis who wishes to seek support on Facebook from friends and family. They find it completely appropriate to share photos of their test results and other medical information on Facebook to keep their friends and family updated. While I believe most people would perhaps share updates on their health states, they would not wish to share their medical records in such detail online. There are, of course, instances where broader norms should trump individual privacy choices even if they are value-centered towards greater ethical aims - especially those that cause obvious harm to others (e.g., “sharenting” (Steinberg, 2017)). While these discussions are necessary, I will not be exploring those conflicts in this thesis.

c Values and Autonomy

Diving into conceptions of autonomy furthers the link between values and privacy decision-making. Being autonomous, or our ability to self-govern, has always been more than choice. While this work is interested in exploring our individual, *personal autonomy* when making privacy decisions, we can see links between autonomy and values within the realm of *moral autonomy*. Immanuel Kant posits that we are rational agents with the capacity to govern our actions according to higher principles (Kant, 1959). In Kant’s case, our ability to reason also suggests that we also have intrinsic rights as persons. From here, he further derives his famous rule for governing our actions: *the Categorical Imperative*.³² The most relevant formulation of this imperative for our exploration of personal autonomy understands humanity as ends in themselves. Simply put, this principle states that we should view people as moral agents with their own desires, aspirations, goals, rights, and inherent dignity, and never as mere means-to-an-end. Not only does this lay the groundwork for the basis of respecting autonomy due to human dignity, but we are understood as agents that are led by “something higher” than impulses or stimuli.

In Kant’s case, this “something higher” were universal principles and maxims. Returning to *personal autonomy*, there is also volitional autonomy, mentioned earlier in discussion of “bright patterns,” where autonomy involves our ability to synergize our “first-order” desires in accordance with our “higher-order” values (Dworkin, 1988b; Frankfurt, 1971). Philosopher Suzy Killmister, who will be discussed in greater detail in the next section, further embraces values as “clusters of commitments” at the core of our deliberations and actions in her theory of autonomy.

Therefore, for a data privacy decision to be an accurate reflection of a user’s values, the decision must be sufficiently unhindered by external forces such as dark patterns. These patterns, as we saw in earlier, greatly diminish user autonomy in notice-and-consent privacy management regimes. They must also be sufficiently engaged with the privacy decision to make a conscious, value-centered decision – no “mindless click throughs.” Autonomy is central to these value-motivated and value-reflective data privacy choices and is central to what a value-centered privacy choice is.

Section 2.4 From Theory to Practice: Designing for Value-Centered Privacy Decisions

The relationships between autonomy, values, and privacy have also intrigued some researchers in the human-computer interaction (HCI) and user experience (UX) communities. These researchers are operationalizing these relationships by *designing* for wellbeing (rather than attention or other factors) and by applying insights from cognitive and behavioral psychology towards more benevolent ends (Cox et al., 2016; Peters et al., 2018; Sandhaus, 2023; Terpstra et al., 2019). In particular, utilizing friction – that is, introducing small obstacles during a user’s interaction with technology – is of particular relevance to this work (Cox et al., 2016). In this case, researchers are hoping friction will induce a shift from “fast” to “slow” user thinking and thereby facilitate more autonomous decision-making. As described by Kahneman (2011), “Fast” System 1 thinking is automatic, mindless, and ripe with the heuristics and biases detailed in Section 2.2. “Slow” System 2 thinking is conscious, deliberate, and mindful. By introducing this friction, researchers are aiming to promote more System 2 thinking and therefore more deliberate

³² Universal law formulation: “Act only according to that maxim by which you can at the same time will that it should become universal law” (Kant, 1959, pg. 39).

interactions with technology. A recent paper further proposed using friction, among other criteria, to re-define and design bright patterns “that prioritizes users’ well-being and goals over their desires and business’ objectives” (Sandhaus, 2023, pg. 3).

Concerning data privacy, friction has been proposed by Terpstra and colleagues (2019) to promote user comprehension of privacy policies and encourage users to make better privacy choices. Notices-as-friction (e.g., “are you SURE you want to share your data?” on a pop-up notice) could be used to encourage reflection, learning, and to draw attention to one’s own underlying beliefs and values (Terpstra et al., 2019). Terpstra and colleagues state that optimizing individual choice is preferable to the soft paternalistic measures such as nudges because a “decision is only truly meaningful when made deliberately.”

Similarly, I wish to understand respecting autonomy in privacy-decisions as designing *for it* - in this case, conceptualizing autonomy in a manner that brings the role of an agent’s personal values to the forefront.³³ I wish to then utilize this understanding to promote more meaningful privacy decisions, ones that would help them act according to their values. In other words, I aim to design *for* value-centered privacy decisions.³⁴

2.4.1 Personalized Privacy Assistants

To design value-centered privacy decisions, I am particularly interested in building off a current technology – personalized privacy assistants (PPAs). PPAs are machine learning assistants that can provide personalized privacy notifications for a user based on their privacy preferences (Das et al., 2018; Liu et al., 2014, 2016; Story et al., 2020; Warberg et al., 2019). Privacy assistants are part of a much larger landscape of related technologies that aim to combat privacy-related challenges, with privacy-enhancing technologies (PETs) being the most well-known. PETs are a collection of techniques and design choices inspired by privacy-by-design,³⁵ such as differential privacy and anonymization techniques (Garrido et al., 2022; Heurix et al., 2015). In contrast, personalized privacy assistants are constructed to help users manage the overwhelming number of privacy notices they face. Their personalized notifications are therefore excellent starting points if we wish to “slow down” users and facilitate more conscious, reflective, yet not overwhelming, value-centered choices.³⁶

2.4.2 Additional Theory, Scope, and Constraints

To design such an assistant, we will firstly require a means of further understanding and conceptualizing what constitutes a value-centered privacy decision. In addition, we will

³³ Exploring ways of promoting autonomy in technology is also a component of VSD (Friedman et al., 2008). My approach – designing for value-centered privacy decisions – however, is different from “classic” VSD in its emphasis on the individual, and in its understanding of the relationship between autonomy, values, and privacy decisions. It could be considered, broadly, VSD, in that it is theoretically grounded in the idea that values are promoted by technological design (see Section 2.3). However, this thesis does not strictly follow the tripartite methodology outlined in Friedman et al. (2008).

³⁴ One might rightly point out that designing for value-centered privacy decisions is a “nudge” – it is, but more in the spirit of a “bright pattern” proposed by Sandhaus (2023) than the ones defined in Grassl et al. (2021). The power is with the nudgee, not the nudger; it is a means of helping us follow through with that desire in our data privacy choices. Others have also argued that such nudges are consistent with autonomy (Killmister, 2017; Schmidt & Engelen, 2020).

³⁵ Privacy-by-design are a series of design choices that protect the user’s data privacy *by default*. For more information, see: Cavoukian (2009).

³⁶ The rationale behind and details of privacy assistants will be more deeply explored in Chapter 3.

require the means of *operationalizing* values to identify, measure, and assess them for the assistant. Lastly, we also need to define *which* privacy decision point we are aiming to design for. Here, I will: 1.) select a theoretical lens to further conceptualize value-centered privacy decision-making; 2.) define the scope and context this thesis will focus on; and 3.) identify a means for operationalizing values.

a 4DT: A Lens for Value-Centered Privacy Decisions

As the previous sections demonstrate, value-centered reflection and action cannot occur if autonomy is frustrated. I aim to, essentially, design for autonomy—with the goal of providing more meaningful privacy decisions that reflect who someone is and what they value.

We need to, therefore, define and identify conditions for autonomy in data privacy decisions. To accomplish this, we can look to the wealth of autonomy literature in philosophy to identify a value-centered conception of autonomy that fits the context of privacy decision-making.

For this investigation, we will need a theory of autonomy that fits three criteria.

Firstly, the theory must have *personal values* at its core. This is to capture the implicit role of our personal values in our data privacy choices, as well as capture the essence of what a value-centered choice is. As discussed in Section 2.3, a value-centered privacy decision is more than “just choice” – it is normatively linked to our ability to self-govern according to our values.

Secondly, the theory must be *reasonably systematic* and practical to conceptualize, assess, and improve data privacy decisions. It must be granular enough to be able to identify instances where PPA design choices are not best facilitating value-centered choices, as well as point to potential avenues of *designing for* value-centered choices in an assistant.

Thirdly, because we are looking at individual data privacy decisions and personal values, it must be a theory of *personal autonomy*. However, to fully meet the first criterion, it must be more than “just choice” – it must involve self-governance according to one’s values. This interplay between the first and third criterion can be captured in the “apathetic user” phenomenon – the phenomenon, first introduced in Section 2.2, where a user becomes so overwhelmed by privacy notices to the point of apathy. To meet the first and second criterion, we must choose a theory of autonomy that has a normative basis that does not allow for one to “choose” to not take on any self-governing commitments or policies concerning their data privacy choices. This, with all its nuances, is quite tricky to meet. While we are not aiming to assess an agent’s ability to exact an objective moral law upon themselves,³⁷ there must still be some normative incentive to *not* become an “apathetic user” by succumbing to notice fatigue – that is, to remain self-governing.

Suzy Killmister’s Four-Dimensional Theory of Self-Governance (abbreviated hereafter as “4DT”) meets these criteria. 4DT divides autonomy into four dimensions – *self-definition*, *self-realization*, *self-unification*, and *self-constitution* (discussed in greater detail in Chapter 3) (Killmister, 2017). Briefly, the first dimension, *self-definition*, is concerned with personal identity: *self-definition* assesses the level of internal consistency between the goals, beliefs, commitments, and values that make up our personal identity. Our goals and beliefs commit us to be or act a certain way, and similar commitments can be further clustered into values. The second dimension, *self-realization*, is concerned with

³⁷ This is not to imply that privacy decisions cannot have moral weight, but merely that it is not what we are investigating here. For an example of moral duties towards privacy disclosure, Allen (2013) draws on Kant to argue that we have a duty toward ourselves to protect our privacy out of respect for ourselves.

Background

practical agency – that is, our ability to deliberate, form a conclusion, for an intention, and act. *Self-realization* assesses the level to which our conclusion from practically deliberating aligns with our intentions (*internal self-realization*), and the degree to which our actions align with our intentions (*external self-realization*). The final two dimensions concern the relationship between one’s personal identity and practical agency. The third dimension, *self-unification*, concerns whether our actions are consistent with our personal identity. The fourth, *self-constitution*, has both foundational and applied aspects. Foundationally, it concerns our ability and willingness to take on and form commitments – and therefore be autonomous agents deserving of respect. If we take on *no commitments*, we cannot be autonomous – we can’t engage in the process of *self-defining*, *self-realizing*, and *self-unifying*. *Self-constitution* can also be applied to the other dimensions to assess the degree or *quality* of our *self-defining* attitudes, deliberations, and intentions. If an agent takes on commitments, but the commitments are weak, and they do not take seriously the process of deliberating and acting consistently in their lives, their *self-constitution* is less. Definitions for these dimensions are outlined in Table 2-1, and the process of being a self-governing, autonomous agent according to these four dimensions is outlined in Figure 2-1.

To the first criterion, *self-definition* has links to *personal values* – it involves forming commitments on how to be and act in the world which, in turn, are clustered into values. “When an agent values x [...] she is committed to doing or being certain things” (Killmister, 2017, pg. 57). Self-governance according to one’s identity and values is at the core of 4DT. While other theories, such as *principled autonomy* (O’Neill, 2002), also centrally locate self-governance in their theories, these are less concerned with *personal values* than with broader norms and principles.

4DT also meets the second criterion: it categorizes autonomy into four distinct, accessible, and practical dimensions for our investigation of values and privacy decisions. This is the advantage of 4DT over other theories of autonomy that concern autonomy more generally. These theories lack the granularity necessary to identify features and to design a privacy assistant that promotes value-centered choices.

Table 2-1: The dimensions of 4DT

Dimension	Definition
<i>Self-definition</i>	Personal Identity: Consistency between the goals, beliefs, and values that make up our personal identity
<i>Self-realization</i>	Practical Agency: Deliberating, intending, and acting coherently
<i>Self-unification</i>	Consistency between Personal Identity (values) and Practical Agency (actions)
<i>Self-constitution</i>	Foundational: Willingness and ability to take on commitments Applied: Taking seriously developing one’s personal identity (<i>self-defining</i>), utilizing one’s practical agency (<i>self-realization</i>), and ensuring unity between the two (<i>self-unification</i>)

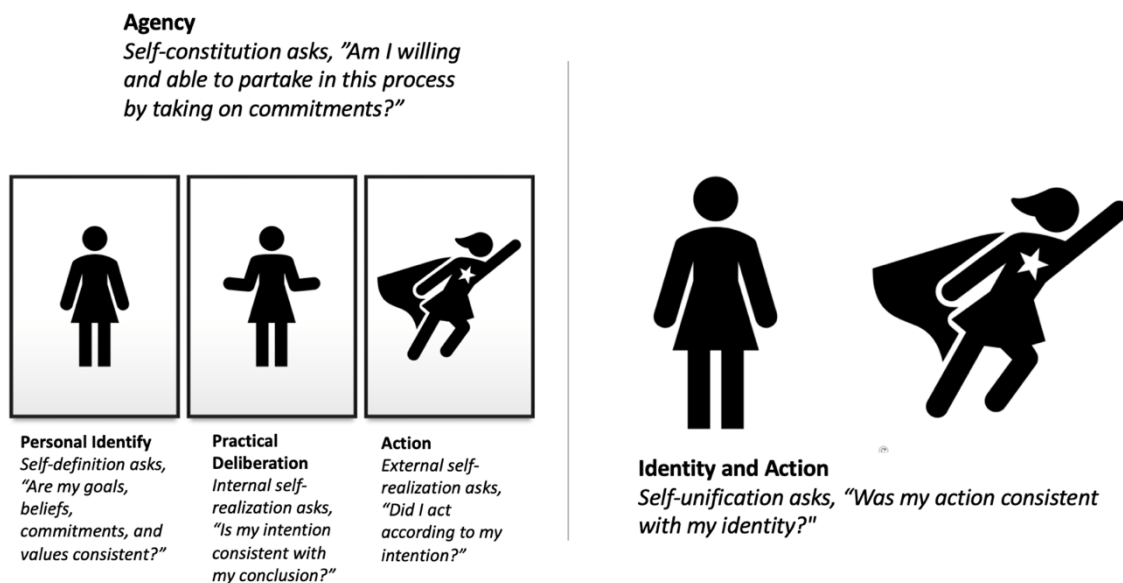


Figure 2-1: Self-governance in action according to 4DT

4DT can also account for the final, and most difficult, requirement – a theory of *personal autonomy* that also does not allow for apathetic choice – that is, one cannot “choose” to not take on any self-governing commitments or policies. 4DT is indeed a theory of personal autonomy, as it is not concerned with whether the ends that are brought about are necessarily the morally correct ones but rather whether the individual is self-governing according to their identity and values. Normativity for 4DT is derived from an agents’ ability to take on or reflect upon their own commitments – that is, we *must* be committed to something we should do or become (Killmister, 2017). These agency requirements are captured in *self-constitution* (Figure 2-1). Thus, someone who refuses to take on any commitments or intentions whatsoever – such as an “apathetic user” – cannot be said to be highly *self-constituting*.

Fitting this last part can be a challenge for closely related theories of personal autonomy that could meet the first two criteria and is, ultimately, where 4DT stands out. Returning to Dworkin, for example, if the “apathetic user” decides to take on no commitments concerning privacy decisions after the highest order of reflection, it could be considered autonomous (Dworkin, 1988a).³⁸ The same would hold true for Frankfurt and his theory of free action, where to act freely – autonomously – we must act in accordance with our second-order volition (Frankfurt, 1971). An “apathetic user” could be considered autonomous if acting in accordance with their highest-order volition or, in other words, desiring to desire not to take on any commitments pertaining to their data privacy.³⁹

Besides meeting these three critical criteria, 4DT also has the added strength of seeing autonomy as something that can *be designed for and promoted (positive freedom)*, rather than purely libertarian notions of leaving the agent alone (*negative freedom*) (Berlin, 1969). In fact, as we will see in Chapter 3, the theory strikes quite a balance between

³⁸ There are, however, instances where someone may act in a manner *similar* to an “apathetic user” while still being *self-constitutive*. For examples, please see Section 2.2 or footnote 39 below.

³⁹ Of course, someone may *genuinely endorse* giving away their data, and such an action would not be considered value-inconsistent with either Killmister or Frankfurt’s theories. The key here is that the “apathetic user” has come to no longer take on any *commitments* about privacy, or, in other words, they forgo being self-governing in this regard. Frankfurt’s theory could accommodate this as autonomous action while 4DT’s *self-constitution* dimension explicitly captures the autonomy-minimizing effect such a policy has on one’s autonomy.

positive and negative freedom.⁴⁰ This is reminiscent of Childress (1990) and original understandings of respecting autonomy from bioethics; we not only have a duty *not* to manipulate (leave the agent alone) but also to actively *promote* autonomy if we wish to truly respect it. To this point, Killmister concluded her book with a discussion of nudges and autonomy (Chapter 9). She states that *enhancing* autonomy could be one means of marrying our growing understanding of human cognitive biases with our duty to respect autonomy:

*“If we’ve been reading our Kant, we may be inclined to think that **our capacities for autonomy command a kind of awe** [...] Without wishing to cast aspersions on the grandeur of human autonomy, I propose that it is far more fruitful— and far more accurate— to see both our status as autonomous agents, and our achievements of local autonomy, **as fragile, precarious, and inevitably dependent on the support of others**. To turn away from the opportunity to aid one another in our autonomy presupposes a robustness and independence that none of us in fact attain. While none of this is incompatible with responding to autonomy with awe, it does suggest that **awe needs to be supplemented with care**. To show respect for a fragile and precarious good requires **nurturing it, and bringing out its potential**, rather than standing aloof and contemplating it from afar. **When we aid one another to be more autonomous, we are showing a consideration for one another’s autonomy that is anything but disrespectful**” (Killmister, 2017, pg. 183, emphasis added).*

This harmonizes well with what we aim to do – *promoting, enhancing, and respecting* autonomy – by designing for value-centered choices.

b Scope: Why Smartphones, Why Apps?

There are many different individual privacy decisions that could be explored and could benefit from a value-centered approach. Each of these different points (e.g., cookie pop-ups, software privacy policies, smartphone privacy permissions) will also very likely involve different value dynamics. It is not feasible, unfortunately, to explore them all. While this thesis will touch upon data privacy decisions more generally, this work will be particularly focused on when users are deciding to download smartphone applications on the App or Google Play Store. There are a few reasons why this privacy decision point is of interest.

Firstly, surveying the literature pertaining to the relationship between values, smartphone, and/or privacy seems to suggest that the three are intertwined (Alashoor et al., 2015; Nurwidiantoro et al., 2022; Obie et al., 2021; Perera et al., 2019; Shams et al.,

⁴⁰ This can also be understood as the ongoing theoretical tension between *relational* autonomy and *personal* autonomy – that is, the ways oppression, power, and social dynamics can limit our choices that may not be captured by a purely individualized approach. Killmister argues that autonomy need not have an explicit relational component, although social relations are indeed relevant to the degree of autonomy we enjoy. For example, if someone or something externally blocks you from doing as you have chosen to, there is a lessening of autonomy (failure of *external self-realization*) Killmister, 2017, Chapter 8 and 9). In addition, *double binds*, to be discussed more in Chapter 3, can be introduced by external, structural, or social pressures (Killmister, 2017, Chapter 8). Critically, though, Killmister states that “the four-dimensional theory is *instrumentally* relational rather than *constitutively* relational [...] the issue is not about how we come to have the personal identities we have, but about what we *do* to be autonomous; and what we do, I have been claiming, makes no essential reference to other agents” (Killmister, pg. 87). This can be viewed as unsatisfactory to some and, while it is out of the scope of this work to wade into a full-on defense of 4DT as “the theory” of autonomy, those interested in these critiques can start with: Mitchell-Yellin (2018).

Background

2023). To consider two works, Nurwidyantoro et al. (2022) found that app value statements – for example, stressing the value of privacy for Signal and Focus – seemed to influence the values they identified when exploring relevant values for apps on GitHub discussion forms. In addition, possible links between values and level of privacy concern were proposed by Alashoor et al. (2015). This suggests not only that company value statements can be influential when looking to choose value-consistent apps, but that our personal values, privacy preferences, and app choices are all interrelated in some manner. In addition, like values and privacy, the added relationship between apps and values makes intuitive sense. Considering the Twitter⁴¹ app versus, say, the Bank of America app, it is obvious that different values and considerations would be relevant, or at the very least, given more weight (e.g., connection and security, respectively). I therefore understand app choices as an *implicit* privacy decision, tied up in a complex, interrelation of personal values.

Secondly, smartphones and apps are modern-day utilities that permeate all aspect of life. With the amount of time spent on smartphone apps in 2022 estimated at an average of 5 hours a day per person, and likely rising (Wakefield, 2022), focusing value-centered exploration of privacy on apps is likely to be the most impactful. This increased dependence on smartphones is also fueled through an app’s ability to “seduce” users into giving more data away through gamification⁴² and other strategies (Troullinou, 2017), increasing their attractiveness when considering a study of choice, autonomy, and values.

Thirdly, there are several reasons why looking at the decision *whether to download* an application is appealing. Chitkara and colleagues (2017) have suggested that data privacy sharing cannot be managed through per-app privacy permissions alone because 70% of app data access requests are by the same 10 third-party libraries.⁴³ However, there is something to be said about empowering us to manage data collection in accordance with our values by preventing our data from being collected in the first place. Targeting *the decision to download* could account for this third-party sharing issue by preventing value inconsistent data sharing to a library in the first place. It also provides an extra layer of control over our personal data that could further complement – rather than contradict – other initiatives. Improving control in a manner that increases value reflection at the decision whether to download an application could also increase market pressure for improved products that better align with our values.⁴⁴ In this manner, value-centered application selection could serve as the first line of defense against value-inconsistent data practice and encourage more consistent applications to be made in the future.

⁴¹ During writing, Twitter became X. See: <https://www.theguardian.com/technology/2023/jul/24/elon-musk-reveals-the-new-twitter-logo-x>. I’ve decided to keep using Twitter throughout this work, as this remains the most recognizable.

⁴² Gamification, using game-like elements in non-game contexts. For a summary of the ethical debate around gamification, see: Kim & Werbach (2016).

⁴³ Instead, Chitkara et al. (2017) created ProtectMyPrivacy – an application that provides users with the ability to control which libraries as well as apps have access to their data. The issue of cross-app tracking and third-party libraries also lead Apple to release the “App Tracking Transparency” feature on iOS 14.5 (Apple, 2021). This feature requires smartphone apps to get permission from users to track their activity across many applications.

⁴⁴ Susser (2019) makes a similar argument, although on the level of norms rather than personal values. He argues that privacy notices that disclose a company’s data practices – even if incomplete – could encourage them to meet social norms concerning privacy.

c Operationalizing Values: The Theory of Basic Human Values

To study and operationalize values into a privacy assistant, we also needed the means of identifying, quantifying, and translating values. To achieve this, we can turn to an established theory and methodology of personal human values from cross-cultural psychology: the Theory of Basic Human Values (TBHV).

The TBHV was developed by cross-cultural and social psychologist Shalom H. Schwartz to establish the theoretical basis for measuring *universal* human values across cultures and societies (Schwartz, 1992, 1994; Schwartz & Bilsky, 1987). In this theory, values “(1) are concepts or beliefs, (2) pertain to desirable end states or behaviors, (3) transcend specific situations, (4) guide selection or evaluation of behavior and events, and (5) are ordered by relative importance” (Schwartz, 1992, pg. 4). The *content of the value* is motivated by one or multiple of three universal, human requirements: to meet our individual needs as biological organisms; to coordinate social interaction; or to ensure the welfare of social groups (Schwartz, 1992, 2012; Schwartz & Bilsky, 1987). From the TBHV perspective, values further serve as standards by which we evaluate what is good and what is bad. The result of this evaluation is linked with our emotional responses. A value that becomes “activated” (e.g., threatened or upheld) in an applicable context is closely linked to our emotional response to that situation (Maio, 2010; Schwartz, 2012).

Based on this understanding, Schwartz theorized ten broad, universal human values, constituting a “continuum of motivations” with reasonable predictive power for cross-cultural value analysis (Schwartz 1992, pg. 45-6). These values were originally represented in a circular structure, where those primarily concerned with individual interests (*Power, Achievement, Hedonism, Stimulation, and Self-Direction*) are opposed to those concerned with more collective interests (*Benevolence, Tradition, and Conformity*) (Schwartz, 1992; Schwartz & Bilsky, 1987). He also proposed that *Security* and *Universalism* could serve both individual and collective interests and were thus placed at the boundary between the opposing individualist/collectivist camps.

Empirical analysis, however, suggested that the relationship between values was not circular, but rather, quasi-circular. To test their ten values, Schwartz and colleagues (1992) designed a survey of 56⁴⁵ terminal and instrumental⁴⁶ sub-values based on the ten theorized broad values. This survey, called the Schwartz Value Survey (SVS), was deployed with participants⁴⁷ from 20 different countries and empirically evaluated using Guttman-Lingoes Smallest Space Analysis (SSA) (Guttman, 1968).⁴⁸ This analysis can be used to generate an image, where empirical relations between values are represented by the space between them. In this case, Schwartz and colleagues observed a quasi-circular arrangement of the ten values (Figure 2-2). They also further grouped the ten values along

⁴⁵ Two years after the first 56 survey in 1992, it was updated to a 57-item survey (Schwartz, 1994). The underlying theory, however, remains the same.

⁴⁶ This was based on the work of Milton Rokeach, who suggested that there were two different ways to represent values: those that allow us to obtain desirable ends (“ends values”) or values that describe a valuable way of acting (“means values”). Different wordings of these values are distinguishable by different impacts on our behavior, even if they are representing the same value. This is reflected on Rokeach’s Value Survey, which contains 18 instrumental and 18 terminal values (Rokeach, 1973).

⁴⁷ Participants were either grade-school teachers or, when not available, university students. Researchers aimed to get ~200 participants in each country.

⁴⁸ One can reasonably wonder whether asking someone about what they value generates responses of what they individually value or instead mirrors values that are desirable in their culture. In response to this, Schwartz (1992) states that “if responses were determined by cultural ideals, then we would expect high group consensus with regard to the importance of each value. However, in every sample studied, there was substantial individual variance in response to every single value” (pg. 50).

Background

two dimensions, which corresponded to their similar motivational goal and position on the circle: 1.) Openness to Change (values: *Stimulation, Self-Direction, Hedonism*) vs. Conservation (values: *Security, Conformity, Tradition*), capturing the tension between our desire for individual independence and for order; and 2.) Self-Enhancement (values: *Power, Achievement, Hedonism*) vs. Self-Transcendence (values: *Universalism, Benevolence*), which captures the tension between our own interests and our concern for the welfare of others. Notably, *Hedonism* has aspects of both Openness to Change and Self-Enhancement. Intriguingly, nearly all the values were observed in their studies in the countries investigated, suggesting cross-cultural validity.⁴⁹




Figure 2-2: Schwartz values in a quasi-circular arrangement. Closely related values are in the same color.

Modified from: Schwartz (1992), with value definitions from Lindeman and Verkasalo (2010).

The TBHV and the SVS has since been utilized by social scientists to measure values across different cultures and contexts (Schwartz, 1992, 1994).⁵⁰ Of most relevance to this work is the use of the TBHV as the theoretical basis for analyzing the role of values in smartphone applications. Such work includes: identifying instances of value violation expressed in app reviews (Obie et al., 2021); elucidating user values from GitHub issue reports (Nurwidyantoro et al., 2022); and identifying relevant values in vulnerable communities for app development (Shams et al., 2023). In addition, Schwartz's value theory has also been utilized to explore values relevant in the privacy field. It was utilized to map GDPR rights and principles onto human values (Perera et al., 2019) and propose possible links between values and level of privacy concern (Alashoor et al., 2015). In this work, we will be extending the use of the TBHV to identify values relevant to privacy decision-making to promoting more value-centered choice (described more in Section 4.1).

⁴⁹ *Power, Achievement, and Tradition* were observed in all countries, and all other values were observed in 90%+ of cases (Schwartz, 1992, pg. 39).

⁵⁰ For a summary, see Schwartz (2012).

Section 2.5 Chapter Summary

In this chapter, we have explored the rationale behind and challenges of privacy self-management – cognitive biases, heuristics, nudges, and (bright and dark) design patterns. We have also looked at how we can respect autonomy in data privacy by promoting *value-centered privacy decisions*. To promote value-centered choices, we must first understand the relationship between values and privacy – that is, *how values are involved in data privacy decisions*. I presented two theories to that will help us conceptualize and operationalize values in data privacy decisions – the Four-Dimensional Theory of Self-Governance (4DT) and the Theory of Basic Human Values (TBHV) (Killmister, 2017; Schwartz, 2012). Both 4DT and the TBHV will allow us to design a value-centered privacy assistant – an assistant that builds upon personalize privacy assistant technology (PPAs). In the next chapter, I will conceptualize what constitutes a value-centered privacy decision using 4DT. I will then use this 4DT-understanding of value-centered privacy decisions to understand why we may not always act according to our values and evaluate PPAs. Lastly, I utilize these insights to design a value-centered privacy assistant (VcPA).

Chapter 3 Crafting the Value-Centered Approach to Privacy Decisions

*It doesn't matter what I say
So long as I sing with inflection
That makes you feel I'll convey
Some inner truth or vast reflection
But I've said nothing so far
And I can keep it up for as long as it takes
And it don't matter who you are
If I'm doing my job, it's your resolve that breaks*

Blues Traveler (“Hook”)

Section 3.1 Chapter Overview

As we have seen, privacy notices are not effective at eliciting informed consent for data collection. Here, I explore how they could instead be utilized to create the space for us to make autonomous, value-centered privacy decisions – designing for autonomy, rather than for privacy, in privacy decision-making.

To do this, I firstly conceptualize value-centered privacy decisions using the Four-Dimensional Theory of Self-Governance (4DT) (Killmister, 2017). After applying 4DT to privacy decisions to conceptualize value-centered choice, I identify and define three major areas of autonomy frustration: *notice fatigue*, *lack of relevant controls*, and *nudges*. I then utilize this understanding to evaluate a current system, personalized privacy assistants (PPAs) to inform the design of a value-centered privacy assistant (VcPA). In particular, I consider designing a VcPA for one context – choosing and downloading smartphone apps – and further demonstrate this system using user scenarios. Based on this analysis, I identify three features a VcPA must have to facilitate value-centered privacy decisions: selective notifications, an exploratory process, and suggesting alternative applications.

3.1.1 Collaborator Contributions

The ideas described in this chapter are my (the PhD candidate’s) work. Feedback was provided by PhD supervisors Dr. Heike Felzmann, Prof. Dr. Mathieu d’Aquin, Prof. Dr. Kathryn Cormican, and Dr. Dave Lewis. I would also like to acknowledge the three reviewers at *Digital Society*, who also provided significant, detailed feedback on the ideas presented here.

3.1.2 Relevant Papers and Conference Contributions

Some material in this chapter, including certain text and figures, has been previously published or presented in the following:

- Carter, Sarah E., & Felzmann, Heike (2023). How do we value data privacy? Insights and design implications. To be published in: *Engineering and Value Change* (part of: *Springer Philosophy of Engineering and Technology series*). Abstract available at: <https://zenodo.org/record/8367542>
- Carter, Sarah E., & Felzmann, Heike. (2023, April 21). How do we value data privacy? Initial results from semi-structured interviews. Forum on Philosophy, Engineering, and Technology (fPET2023), Delft, the Netherlands. Zenodo. <https://doi.org/10.5281/zenodo.8204406>
- Carter, S. E. (2022). A value-centered exploration of data privacy and personalized privacy assistants. *Digital Society, 1*(27), 1–24. <https://doi.org/10.1007/s44206-022-00028-w>
- Carter, S. E. (2021). Is downloading this app consistent with my values?: Conceptualizing a value-centered privacy assistant. In D. Dennehy, A. Griva, N. Pouloudi, Y. Dwivedi, I. Pappas, & M. Mäntymäki (Eds.), *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics): Vol. 12896 LNCS* (pp. 285–291). Springer International Publishing. https://doi.org/10.1007/978-3-030-85447-8_25
- Carter, Sarah E. (2021, September 1). Is downloading this app consistent with my values? Conceptualizing a value-centered privacy assistant. The 20th IFIP Conference e-Business, e-Services, and e-Society (I3E2021), Online. Zenodo. <https://doi.org/10.5281/zenodo.8205147>
- Carter, Sarah E. (2021, July 5). A value-centered exploration of data privacy and personalized privacy assistants. CEPE/IACAP Joint Conference 2021: The Philosophy and Ethics of Artificial Intelligence (CEPE/IACAP 2021). Online. Zenodo. <https://doi.org/10.5281/zenodo.8205315>
- Carter, Sarah E. (2021, February 25). Improving notice: the argument for a flexible, multi-value approach to privacy notice design. 30th Annual Association for Practical and Applied Ethics Conference (APPE 2021). Zenodo. <https://doi.org/10.5281/zenodo.8205297>
- Carter, Sarah E. (2020, December 15). Four-Dimensional autonomy in a digital age: Where are privacy notices going wrong? Ends of Autonomy: December Colloquium. Zenodo. <https://doi.org/10.5281/zenodo.8204502>

Section 3.2 Conceptualizing Value-Centered Privacy Decisions

3.2.1 Applying the Four-Dimensional Theory of Self-Governance to Data Privacy Decisions

In this section, we will explore our selected value-centered understanding of autonomy – the Four-Dimensional Theory of Self-Governance (4DT) – in more depth. We will apply its dimensions in the context of data privacy decisions and utilize it to further conceptualize value-centered choice. Recall from Chapter 2 that 4DT consists of four different dimensions. 4DT was selected as the theoretical lens of choice in part because it

categorizes autonomy into these four distinct, accessible, and practical dimensions for our investigation of values and privacy decisions. This is the advantage of 4DT over other theories of autonomy that concern autonomy more generally, which lack the granularity necessary to identify specific features and to design a privacy assistant that promotes value-centered choices. Conceptualizing value-centered choices using 4DT will further allow us to identify critical features for designing a value-centered privacy assistant (VcPA).

To begin, let us start with the first of these dimensions – *self-definition*. *Self-definition* is concerned with personal identity: it assesses the level of internal consistency between the beliefs, goals, values, and commitments that make up our personal identity. While *self-definition* is not particularly relevant to the *act* of privacy decision-making, it does help us better understand how values are involved and reflected in the privacy decision-making process. As we saw in Chapter 2, the centrality of personal values to 4DT, encompassed by the *self-definition* dimension, also makes it a suitable theory for our exploration of value-centered privacy choices. As outlined in Table 2-1 and Figure 2-1, our *self-defining* attitudes – beliefs and goals – commit us to act a certain way. To demonstrate this, let us consider an example. Perhaps I believe too much screen time is bad for my health. I am therefore committed to spending less time on the computer. 4DT understands values as groups of commitments that are oriented towards some desirable end-state. Perhaps I also have the goal of running a marathon. This commits me to run every day, which could be clustered with my other commitment to minimize computer time to say I value *health*.




Figure 3-1: Four-Dimensional Theory of Self-Governance (4DT) as it pertains to individual data privacy decision-making

When considering *self-definition*, it is important to note the influence of *self-governing policies*. These policies shape our *self-defined* values and can have downstream effects on our privacy preferences. *Self-governing policies* are those that dictate what and how to “believe, plan, and value” (Killmister, 2017, pg. 22). Consider, for example, that I am someone who has a policy that I should carefully protect the boundary between my personal and public life. I come to believe that social media data collection practices threaten this boundary, and therefore commit to carefully monitoring the privacy policies of the social media sites I use. This could cluster with other similar commitments, and, depending on their composition, we could say I value *privacy*, *security*, and/or *control* over my life. Assuming that I can practically deliberate and act in a manner consistent with my values, this *self-governing policy* will ultimately affect my privacy preferences and data privacy choices.

The second dimension, *self-realization*, is concerned with practical agency: *self-realization* assesses the level to which our conclusion from practically deliberating aligns with our intentions (*internal self-realization*), and the degree to which our actions align with our intentions (*external self-realization*). Returning to our example, let's say I am considering downloading a running smartphone app that would like to access my health data. I deliberate and determine based on my values and, because I value *health*,⁵¹ conclude that I best ought to allow it access. I form the intention, or in this case, a *privacy preference*, to allow the app access to my health data. I then act on this and allow it access – that is, make a *privacy decision*. This upholds *self-realization*, as I deliberated, formed an intention, and acted coherently.

The final two dimensions are concerned with the relationship between personal identity and practical agency. Applying them to data privacy decisions also comes with a few important caveats. Firstly, *self-constitution* originally pertained to someone's *overall* willingness to take on commitments. An agent could still be considered highly *self-constituting* if their willingness is present in a *range* of areas, even if this willingness does not extend to *all* circumstances or topics. In this case, I am utilizing 4DT as a tool to analyze autonomy constrained to a specific circumstance – individual data privacy decisions – and am treating *self-constitution* as *the willingness to take on any or new commitments pertaining to data privacy decisions*. Secondly, recall from Chapter 2 that *self-constitution* can also be applied to both the domains of personal identity and practical agency. It encompasses the degree and quality of one's *self-defining* attitudes; quality of their practical deliberation and intention-formation; and the degree to which they can act in a unifying manner over *a period of time*. When considering how *self-constituting* they are with respect to data privacy, this would be a longer-term process and not one that can be determined at any one privacy decision point. We can therefore further amend the aforementioned definition of *self-constitution* to one's *willingness to take on any or new commitments pertaining to data privacy decisions at a single privacy decision-making point*.⁵² This brings me to the third and final caveat: *self-unification* is not only concerned with ensuring that one's actions match their personal identity, but also the extent to which the conclusions/commitments generated during practical deliberation can inform transformation (a change in personal identity). Again, this is tied into *self-defining* attitude formation and is not something that is made at one point in time. Because of this, an analysis of data privacy-decision making in terms of the *self-unification* dimension will focus on *the overall coherence of someone's actions to their personal identity*.⁵³ If we return to the case of downloading a running app because I value my health, *self-unification* is upheld because my action to download the app and give it access to my health data is consistent with valuing *health*. *Self-constitution* is also upheld because I am able and willing to form conclusions and intentions concerning whether to download the app.

3.2.2 Conceptualizing Data Privacy Challenges with a 4DT lens

In the previous section, I have presented ideal examples – that is, when users are making privacy decisions according to their values and in accordance with 4DT. As we have seen

⁵¹ For clarity, values will be italicized throughout this work.

⁵² One may point out that such a statement seems internally contradictory, as Killmister (2017) puts *self-constitution* forward as a type of *global autonomy* with an inherently temporal component. However, *self-constitution* is critical for accounting for the “apathetic user” phenomenon. It provides us with a normative basis that does not allow for one to “choose” to not take on any self-governing commitments or policies concerning their data privacy choices. See Section 2.4.

⁵³ This has implications for how I understand *akrasia* in the next section. See footnote 58.

in Chapter 2, this is not always the case. In this section, I explore why we do not always make value-centered privacy decisions using a 4DT-informed approach to data privacy. I firstly conceptualize three previously identified problematic areas when it comes to privacy decision-making in terms of 4DT's dimensions. These are: (a) notice fatigue; (b) a lack of relevant privacy controls; and (c) nudges (including dark and bright patterns).⁵⁴ I also aim to remain mindful of the insights from behavioral and cognitive psychology introduced in Section 2.2 when conducting this analysis, marrying these critical insights concerning our psychological decision-making limitations with 4DT. In addition, conceptualizing these existing data privacy decision-making challenges in terms of 4DT establishes a baseline for assessing personalized privacy assistants (PPAs) – identifying where they are succeeding at promoting value-centered privacy decisions, and where they could be improved when designing a VcPA.⁵⁵

a Notice Fatigue: A Matter of Degrees

First, we look at notice fatigue – that response to the constant barrage of notices and privacy decisions that leads to unconscious “click-throughs” (Schaub et al., 2015).⁵⁶ From a 4DT perspective, the phenomenon of notice fatigue can be understood as a failure to *self-realize*, *self-unify*, *self-constitute*, or a combination of these (Table 3-1). I've broken these down here in terms of severity, or *degrees* – that is, how many dimensions of 4DT are violated. I will present second- and third-degree notice fatigue first, as the first-degree notice fatigue is easier to understand following these more severe autonomy failures.

(1) Second-Degree Notice Fatigue: Self-Realization and Self-Unification Violations

There are two possible manifestations of second-degree notice fatigue, here defined as failures of *self-realization* and (likely) *self-unification*. They differ in whether the failure involves *internal self-realization* (forming an intention, or privacy preference) or *external self-realization* (acting on one's preference) (Figure 3-1).

We can firstly consider a failure to *external self-realization*. Imagine a user who deliberates, concludes, and *forms the corresponding intention not to share* certain kinds of personal information (e.g., location) because they deeply value *self-determination* and *control* over where their personal information goes. Instead, they exhibit notice fatigue by “just clicking through” all notifications. This is a failure of *external self-realization* because their intention (privacy preferences) and action are not consistent.

We can also imagine second-degree notice fatigue as a failure of *internal self-realization*, where their behavior is the result of their intention to click through despite having concluded they ought not to do otherwise. Using the same example, the user concludes, through practical deliberation, that they best ought not to share their personal information, but nevertheless *intends to share* it; their *privacy preference* is now incoherent with their personal identity. They act on this, clicking through all notices. *External self-realization* would be conserved, as intention and action remain consistent, but *internal self-realization* would be violated.

⁵⁴ These problems were introduced in Chapter 2 and are explored more deeply here by utilizing a 4DT lens.

⁵⁵ While I am not aiming to design for informed consent, it is interesting to note that Killmister describes how the conditions of informed consent can be captured in terms of self-constitution (competence) and self-realization (voluntariness and proper information) (2017, Chapter 7).

⁵⁶ Detailed in Section 2.2.

To understand why *self-unification* is (likely) violated in both cases, we must first introduce two new terms. The first failure of *external self-realization* would constitute what Killmister (2017) refers to as *weakness of will*, and the second one of *internal self-realization* what she would define as *akrasia*.⁵⁷ *Weakness of will* is when we intend to act a certain way – for example, not share our location, but consent to sharing it anyway. *Akrasia*, for our purposes,⁵⁸ is when we form the intention to act against what we practically conclude, based on our values (commitments), we ought to do. This distinction is important for *self-unification*, where the *weakness of will* case always results in a failure of *self-unification* while the *akrasia* does not necessarily result in a failure of *self-unification*.⁵⁹ This is because the “click through” actions that this user takes when fatigued will be inconsistent with their personal identity and their values. There is only one situation where this would not be true, and this would be the “lucky akratic” (Killmister, 2017, pg. 64). The “lucky akratic” forms an intention against their values and yet, unbeknownst to them, actually acts in accordance with them. To demonstrate this further, consider a user who is considering downloading WhatsApp at the recommendation of their friends. They deliberate and realize that they are not comfortable with having any data, even metadata,⁶⁰ shared to WhatsApp’s parent company, Meta. They instead intend to download the app and go download it. However, at the time of download, they forget the name of the app, and they download Signal instead of WhatsApp. Lucky for them, Signal is more privacy preserving than WhatsApp and does not collect or store user metadata.

As this example demonstrates, the “lucky akratic” in our case of privacy decision-making is likely pretty rare. It is unlikely that someone takes the effort to deliberate on WhatsApp and then “forget” at the time of download the name of the app. In addition, I would argue that the manifestation of second-degree notice fatigue that is occurring *most* of the time is one of *external self-realization* (*weakness of will*), rather than *internal self-realization* (*akrasia*). While I cannot claim to know the internal states of users, I think most scholars would agree that the major challenge of notice fatigue is that we are not following through on our privacy preferences when faced with, simply, an incredible barrage of

⁵⁷ Some readers may be used to using *weakness of will* and *akrasia* interchangeably to mean “any action we take against what we best ought to do.” Killmister specifically breaks such actions into two types based on whether the autonomy failure is occurring *internal* or *external self-realization* failures is relevant. Also note that in the case of the “lucky akratic,” who may appear to be acting in a coherent manner while still being *akratic* according to 4DT.

⁵⁸ Those who have done a deep reading of Killmister (2017) will note that I have simplified her definition of *akrasia* here. She also outlines a more complex case of *akrasia* where an agent deliberates and *brackets* an aspect of their personal identity, not including it as a consideration in their practical deliberation. This results in an *akrasia* failure when the agent is deliberating, rather than when the agent is forming their intention. *Bracketing* is the result of a complex interplay of different aspects of one’s personal identity (e.g., if one’s religious beliefs require them to *bracket* their sexual desires), and it results in actions that are externally self-realizing but not *internally self-realizing* nor *self-unifying*. This violation of *self-unification* can occur at: 1.) the action taken; and 2.) the agent’s *akrasia* at the deliberation stage. As stated in Section 2.2.1, I am utilizing *self-unification* in the context of privacy-decisions to mean the coherence between one’s actions and their personal identity. This does not mean that we do not *bracket* when making privacy decisions (we likely do). However, for my purposes and aim of promoting more value-centered privacy decision-making, the impact both forms of *akrasia* have are the same – 1.) the formation of a less-than-genuine privacy preference, and 2.) the corresponding incoherent action. I therefore use *akrasia* throughout this work in terms of what Killmister calls “type 1” *akrasia* (failure at intention forming), rather than the “type 2” *akrasia* (failure from *bracketing*). For those interested in reading more on the different types of *akrasia*, see Killmister (2017, pg. 36-39).

⁵⁹ See footnote 58.

⁶⁰ While WhatsApp does not share message content with Meta, it does share metadata. This metadata includes “when [your message] was sent and your IP address” and “X is on Y’s phone and they are messaging each other every evening at around 8pm for an hour” (O’Flaherty, 2021).

notices.⁶¹ It takes someone exhibiting great willpower to do so. In addition, users *have documented intentions and preferences* and seem to care about their privacy actions - they are just not *following through* and acting according to them (Norberg et al., 2007).⁶²

(2) *Third-Degree Notice Fatigue: A Failure of Self-Constitution*

The highest degree of notice fatigue – third-degree notice fatigue – involves a failure to *self-constitute*. As mentioned in Section 2.2 a user also runs the risks of becoming an “apathetic user” – that is, becoming so overwhelmed by privacy notices that they stop forming new intentions and commitments about their data privacy. In this case, not only are they acting contrary to their values (failure of *self-unification*), but they are not willing to take on commitments, practically deliberate, and form intentions (all failures of *self-constitution*). Indeed, *self-realization* becomes no longer analyzable because the action becomes, essentially, “mindless.” They just “click through” notices. This is considered more severe because *self-constitution* – taking on commitments – is the bare minimum of what is required to be an autonomous agent. Taking on no commitments at all is the equivalent of staring at a rock-climbing wall and refusing to try to climb. Those who become fatigued part way up the wall and decided to climb down at least gave it a shot.

It is important when discussing third-degree notice fatigue to clearly define what exactly is meant by an “apathetic user.” As presented in Chapter 2, the term “apathetic user” used here is meant to capture those who would prefer to be more data protective but feel overwhelmed by notices to the point of “no longer caring,” or apathy. The purpose of this term is to capture a state in which an individual 1.) has internalized apathy and 2.) takes value-inconsistent action. From a 4DT lens, these are failures of *self-constitution* (internalized apathy) and *self-unification* (inconsistent action).⁶³ This, of course, comes with some nuance. For example, someone who reflects upon their data privacy choices and decides to “click-through” all privacy notices because they, say, believe that sharing their data will stimulate technological progress and/or to improve their user experience, are not meant to be captured as “apathetic users.”⁶⁴

⁶¹ This could also be due to a *double bind*, described in (3).

⁶² One might counter that by probing (asking) users about their privacy preferences, we are encouraging participants to have preferences. There are reasons that suggest that this is the wrong way to look at it. By asking users about their privacy preferences, we are, at best, bringing to light their pre-formed self-defined attitudes towards privacy, or, at worst, stimulating a point of reflection where they self-constitute and take on commitments towards privacy. People are largely willing and able to form opinions about privacy. The only exception would be the “apathetic user,” whom I touch on in (2).

⁶³ It could be argued that such actions constitute “digital resignation” as defined by Draper & Turow (2019), and not apathy. If this is the case, such actions could be seen as completely reasonable responses given the sense of futility in the face of such corporate power and ineffective privacy self-management regimes. In my view, the difference between the “apathetic user” and someone experiencing “digital resignation” is that someone who is resigned *still cares about privacy* but feels like their actions are ineffectual, while someone who is an “apathetic user” has, over time, come to no longer take on any commitments about privacy – that is, *no longer care about privacy*. This distinction can also be seen in Draper & Turow (2019, pg. 1834), who note that those who “are resigned can exhibit behaviors that are similar to those who express indifference about digital surveillance, [and] resignation can obscure signals that people care deeply about privacy.” The “indifferent” individuals in this case are (likely) “apathetic users” (see footnote 64 below).

⁶⁴ However, I would argue that the instances where someone *genuinely wills* that are all their data be collected are likely rare. Few of us “click through” notices because we believe in what we are doing, and more so because we feel overwhelmed.

(3) *First-Degree Notice Fatigue: Double Binds*

First-degree notice fatigue is interesting because, unlike the others, it is a kind of notice fatigue that does not have to do with becoming overwhelmed or demotivated. It instead has to do with *actively choosing not to become fatigued*. This relates to what the agent values. Many users, for a variety of reasons, value *efficiency*. We can, for example, imagine a user who holds this value in equal esteem to *control*. In the current privacy notice landscape of an overwhelming barrage of privacy choices, they will not be able to *self-unify*. No matter how they act, they will act contrary to their values – if they choose to (try to) engage with *every* privacy notice that comes their way, they will be violating their value of *efficiency*; if they take the alternative action and choose to engage with notices in a limited manner in the name of *efficiency*, they will violate their value of *control*. They will be in what Killmister refers to as a *double bind*⁶⁵ – a situation in which, no matter how they act, they will be acting against a value that defines their personal identity. We can understand one’s failure to follow through not from *weakness of will* (second-degree) or apathy (third-degree), but because the sheer number of privacy notices means they simply *cannot* act in a manner that honors both their values.

In contrast, we can consider a similar agent who also values *efficiency* and *control*. However, unlike our previous example, they do not hold *efficiency* and *control* in equal esteem – they value *control* less than *efficiency*. While the current landscape of privacy notices means that the two will be inevitably forced into tension with one another, the situation is still *resolvable* in terms of 4DT. This agent deliberates, decides they have most reason to *not* engage with most of the privacy notices they receive, and act accordingly. While this agent’s actions are identical to that of a someone exhibiting notice fatigue (regardless of degree), they are fully autonomous and therefore would be making a value-centered privacy choice.

b Lack of (Relevant) Controls

There are also challenges at the level of the privacy controls themselves. There are a plethora of different privacy notice designs with different degrees of privacy control (Utz et al., 2019), which are not always satisfying for the user (Felt et al., 2012). If a user forms an intention (privacy preference) but the control is not present, this would be a frustration of *self-realization*. Because they are not acting according to their values, this would also constitute a failure of *self-unification* (Table 3-1).

A lack of relevant controls can also be understood in the context of smartphone apps as the lack of availability of an alternative app of similar value. In this case, we can use the concept of *double binds* – although, unlike those present in first-degree notice fatigue, these involve other values outside of *efficiency* (Table 3-1). We can consider the example of social media and messaging applications. We can imagine a user deliberating whether to download a social media app such as Instagram. While they greatly value *social connection*, they also value *control* over their life and data equally. They must decide whether to join social media and allow their data to be collected or to abstain, neither of

⁶⁵ The concept of *double binds* comes from Part II: Applications of A Four-Dimensional Theory of Self-Governance of Killmister (2017), where she explores, among other aspects, the autonomy of the oppressed, and how they are often forced into *double bind* situations. For our purposes, *double binds* are understood here to be when design or structural decisions force two values that an individual holds in high esteem into irreconcilable tension with each other. This is different from value tensions, which can be resolved through weighing one’s values and practical deliberation. We will see examples of both of these situations in Chapter 5.

which will be in full accordance with their values. This will only be a *double bind* if an alternative with similar social value to them is not available. This would occur if most of their friends and family are on Instagram and Instagram is the dominant service in their country or community.⁶⁶ As a second example, we can also return to the hypothetical user from earlier in this chapter who values *health* and downloads an app to help track their runs. Let's say that this app also *requires* access to *all* your health data on your phone (e.g., iPhone's "Health").⁶⁷ We can also consider that this user values something else equally – say, *security* – which means that they ought not to allow an app to access their sensitive health data. This puts them in a *double bind* – no matter how they act, they will violate their values. If they download the app, they violate their value of *security*; if they do not, they violate their value of *health*.

c Nudges: Bright, Dark, or Somewhere In-Between?

We also need to consider the (dark or bright) patterns utilized on privacy notices from a 4DT lens. As mentioned in Chapter 2, there are a host of cognitive heuristics, biases, and design strategies that can be used to nudge one to give away their data (dark) or to take more privacy-preserving choices (bright). Both can be problematic from the standpoint of respecting autonomy.

Besides re-conceptualizing these problems through the lens of 4DT, it is also critical we understand *when* nudges and (soft) paternalism are acceptable in order to inform our design of value-centered choices. In Section 2.4, I mentioned how I want to design for value-centered choice in a manner similar to that of Sandhaus (2023) – that is, reclaiming the term “bright pattern” as those in line with our values and interests. 4DT helps us define the boundaries for kinds of bright patterns - nudges and paternalism from this approach are more nuanced, not all-or-nothing. Indeed, Killmister (2017) thoroughly explores the impact of nudging on autonomy in Chapter 9, and a brief exploration of the main points relevant to this thesis are described below.

(1) Using Nudges Wisely: Self-Binding

From a 4DT perspective, nudging in data privacy decisions is appropriate in two instances. Each of these cases constitute a form of *self-binding* – here used to mean a method of helping oneself follow through with what is consistent with one's personal identity (value and commitments), thereby upholding *self-unification*. For it to constitute *self-binding*, the agent must have enacted the *self-binding* method (in this case, nudges) *willingly*.⁶⁸ To understand why this must be so, consider a person who has the commitment *not* to have their deliberation and actions affected by nudges.⁶⁹ We could say that this commitment is part of the larger value *self-determination*. Trying to promote more value-centered privacy choices using a nudge would inherently violate their values. Therefore, a system that aims

⁶⁶ See the example of Shauna in Table 3-1 and further explanation of social media/messaging *double bind* cases in footnote 70.

⁶⁷ Apple has two levels of control when it comes to the app accessing health data. They have “allow app to write data,” which means that app can add to your health data (say, the length or your workout). They also have “allow app to read and write data” which means they can also *access data written by other apps*. In this example, this app *requires* the latter. See: <https://support.apple.com/en-ie/guide/security/sec88be9900f/web>

⁶⁸ Killmister (2017) notes that nudges cannot, generally, compensate for *weakness of will* – except, I argue, in cases of *self-binding*. See footnote 81.

⁶⁹ This example is adapted from Killmister, 2017, pg. 172.

to promote value-centered privacy decision (such as the VcPA, to be explored 2.3) must also be entered into as a form of *self-binding* – in simpler terms, it must *always* be optional.

Table 3-1: Understanding data privacy challenges through the lens of 4DT

Challenges	Relevant Dimensions	Relevant Concepts	Example
Notice Fatigue	<i>Self-constitution</i> (3°) <i>Self-unification</i> (1°/2°) <i>Self-realization</i> (2°)	The Apathetic User Double Binds Akrasia Lucky Akratic Weakness of Will	Abdul values his <i>personal security</i> and intends to not share his location information, but “clicks through” location sharing requests when downloading apps to his smartphone. He is exhibiting <i>weakness of will</i> .
Lack of Relevant Controls	<i>Self-realization</i> <i>Self-unification</i>	Double Binds	Shauna wants to download WhatsApp because she values <i>social connection</i> and wants to connect with her friends. However, she also values <i>trustworthiness</i> and is concerned about the metadata shared with parent company Meta, who she finds untrustworthy. No matter how she acts, she will violate an aspect of her personal identity (either her value of <i>social connectivity</i> and <i>trustworthiness</i>). Shauna is in a <i>double bind</i> . ⁷⁰
Nudges	<i>Self-realization</i> <i>Self-unification</i>	“Blanket” Nudges Akrasia Weakness of Will Non-Deliberative Agent	Ash values <i>innovation</i> and believes in sharing their data with companies to improve the quality of services provided. However, they receive a “blanket” nudge that discourages them from sharing their data.

This noted, let us turn to the first instance of appropriate notice use. Firstly, a nudge is appropriate if it helps a *deliberative* agent follow through on what they intend to do (*external self-realization*). In these instances, it acts as a deterrent to the *weakness of will* phenomenon – when we intend to do something, and instead act another way.⁷¹ This also will help fulfill *self-unification*, as the agent will be taking an action consistent with their personal identity. In the second instance, cases of a *non-deliberative agent*, nudges can also

⁷⁰ This case is not as extreme as cases of oppression in which the term *double bind* is usually used in the literature (see footnote 65). However, WhatsApp is a dominant messaging service in many countries, with as < 90% of messaging users in Brazil, India, Germany, and the Netherlands, to name a few (*WhatsApp Penetration Rate among Global Messaging App Users as of April 2022, by Country*, 2023). In such an environment, WhatsApp has become a near essential means of modern social interaction. If we assume that Shauna is in one of these countries and has a high valuing of *trustworthiness* and *social connectivity*, we would consider this a *double bind*. This does not mean that *everyone* who values *trustworthiness* and *social connectivity* will be in a *double bind* when considering WhatsApp, even if they are in a WhatsApp-dominated country. Some in these situations will still resolve the tension between these two values because they either value *trustworthiness* or *social connectivity* higher.

⁷¹ This is only true if the agent has chosen to have the nudge as a form of *self-binding*. See Section 3.2.

help us act consistently with our values and personal identities (*self-unification*). In these cases, nudges assist agents by *replacing* deliberation (*self-realization*) when agents are acting in a *non-deliberative* manner (Killmister, 2017, pgs. 63, 65, and 170). Critically, in 4DT, respecting autonomy is not just about letting people do what they want – especially when not acting with intention.⁷² Killmister (2017, pgs. 123-127) uses John Stuart Mill’s bridge example, where stopping someone from walking onto a bridge that is about to collapse is viewed as an action we *ought* to take if we want to respect that individual’s autonomy. For Killmister, intervening in this matter does not violate their autonomy, as there needs to be some consideration, some self-deliberation and intention (*self-realizing*), to be autonomy-violating. We also *must stop them* in such situations if we wish to respect their autonomy.

*“The individual poised to inadvertently cross the perilous bridge is an autonomous agent, and hence has a right to autonomy. **We do not respect that status, though, if we simply let her cross the bridge.** Similarly, we do not respect the autonomy of an individual with a fatal allergy to peanuts if we stand idly by while she orders a meal that we, but not she, know to contain peanuts. It thus turns out to be **too simplistic to say that respect for autonomy means letting autonomous agents do as they have chosen;** there must be something sufficiently autonomous about the actions that the agent is poised to perform, if standing back and allowing her to continue is to count as respecting her right to autonomy”* (Killmister, 2017, pg.126, emphasis added).

Similarly, letting people click “mindlessly” through privacy notices can also be seen as such an issue, where a well-placed nudge could help them *self-unify*. Frequently during privacy decision-making, users are in a task-focused mindset, browsing the internet seeking information or engaging with an app that provides them with some function. In addition, as demonstrated with notice fatigue, there are also far too many choices for us to make conscious, deliberative choices about *all* of them.⁷³ Users are unintentionally walking off data-privacy bridges they may not know are inconsistent with what they value because they are focused on their tasks or something else and are not mindful of the interaction.

(2) *Failures: Blanket Nudges and Akrasia*

However, there are instances where nudges can violate *self-realization* and *self-unification*. There are two such scenarios relevant to data privacy decision-making – “blanket” nudges and frustrating an *akratic* agent.

Firstly, there are “blanket” nudges. In our case, we can consider them “blanket” privacy notices – that is, nudging indiscriminately either to or to not disclose data. As Killmister writes:

*“If nudges are going to augment autonomy, then, **they have to nudge the agent towards actions that accord with their personal identities.** Insofar as nudges are utilized as broad public policies, however, **this outcome cannot be guaranteed.** Unless we assume that all agents share certain goals and values, any given nudge is likely to push some agents towards actions that conflict with their personal identities”* (Killmister, 2017, pg. 171, emphasis added).

⁷² This is related to the idea that respecting autonomy, according to 4DT, also means *promoting* it. See Section 2.4.

⁷³ It would take an estimated 244 hours a year to read all the privacy policies we consent to (McDonald & Cranor, 2008). See: Section 2.2.

From a 4DT, lens, then, dark patterns and (classic, privacy preserving) bright patterns are not autonomous because the pattern may: 1.) encourage the user to act against their intention (*self-realization* failure); or 2.) act in a manner inconsistent with their personal identity and values (*self-unification* failure). Whether these nudges take the form of highlighting “select all” cookies or the broad use of friction in the form the notice itself (notice-as-friction), if it is not individualized, some users will be encouraged to act and indeed act inconsistently with their personal identities. In these cases, the *person designing the nudges decides when and what* the user does, rather than the *agent’s personal identity* (their commitments and values). This is especially problematic when nudges are designed in a more personalized manner to encourage a specific person or group of people to take a certain action by using the cognitive “tricks” that will affect them most.⁷⁴ In summary, indiscriminate, “blanket” uses of nudge are not acceptable.⁷⁵

Secondly, different paternalistic interventions frustrate *self-realization* if an *akratic* user intends to act against their values and personal identity. This point is demonstrated by Killmister (2017, Chapter 9) by using the examples of three vegetarians: Becky, Bob (and Bob*), and Billy. I will summarize the main points here for our purposes with the example of Bristol. Let’s say Bristol wants to be vegetarian and is going to a friend’s house for a barbeque. When she gets to the BBQ, she surveys the burger options available and *akratically* decides to eat a meat burger. She is then presented, perhaps by a well-intentioned friend, with a tofu burger instead, which she then eats and is none the wiser. Even though Bristol’s *self-unification* is improved, her *external self-realization* is violated on top of her *internal self-realization*. Killmister writes that:

“If intervention is to be judged purely in terms of whether or not it enhances the autonomy of the intervened upon agent, then it cannot typically be used to compensate for failures of autonomy caused by akrasia. In most cases, we cannot render the akratic agent more autonomous through intervention— whether that’s surreptitious manipulation of the act she is poised to perform, or physical intervention to prevent the akratic act. Such interventions fail to have any impact on the aspects of autonomy that are causing the problem, and simultaneously frustrate those aspects of autonomy that would otherwise have survived the akrasia” (Killmister, 2017, pg. 169, emphasis added).

Instead, we could try to reason with the *akratic* Bristol, to encourage her to reflect (deliberate) and form new intentions and actions in coherence with their values - even if us raising the issue may not be well-received.⁷⁶ This would further promote their autonomy by giving her an opportunity to deliberate and form new (value-consistent) intentions and take the appropriate *self-unifying* action.⁷⁷

In the case of data privacy, again let’s consider Bristol, who we will say has *akratic* tendencies. She “clicks through” privacy notices (second-degree notice fatigue) despite forming the conclusion that she best *ought not to*, based on her values. Introducing a

⁷⁴ As the case of “hypernudges” – nudges designed specifically for us (Yeung, 2017).

⁷⁵ Unless we can assume the group of individuals share the same value set.

⁷⁶ There is an important catch here – Killmister (2017) notes in Section 3 of Chapter 9 that reasoning with an *akratic* agent can be disrespectful. This is because doing so implies that they lack some reasoning competencies and undermines the intention they made which, while *akratic*, still is deserving of *some* respect because it is made by an *autonomous agent*. I agree with this, which is why this “reasoning” should be the choice of the person to have – a way of following through on their commitments and values. See Section 3.2 for a discussion of this *self-binding*.

⁷⁷ It is also interesting to note that if Bristol *had not* deliberated, but merely took the burger her friend offered her, her friend’s “nudge” (in this case, a default) would be completely appropriate. See the discussion of the *non-deliberative agent* in the previous paragraphs.

classic nudge to encourage her to act according to her values *could* frustrate her *external self-realization* unless proper rationale is given to reason with her. We could consider, for example, the case of using friction to encourage reflection on one's values when making a privacy decision. As mentioned in greater detail in Section 2.4, notice-as-friction could be used to encourage reflection, learning, and to draw attention to one's own underlying beliefs and values (Terpstra et al., 2019). Having pop-up notices that aim to “reason” with someone making an *akratic* data privacy decision could be one option. It is also worth noting that I have claimed when discussing second-degree notice fatigue that any *self-realization* failure by a *deliberative* agent when making a data privacy decision is *likely* happening at *external self-realization (weakness of will)*, rather than *internal self-realization (akrasia)*.

In summary, nudges can be used in a manner consistent with 4DT when consistent with one's personal identity (values) and when such nudges are entered into willingly as a form of *self-binding*. In data privacy decisions, it can also be harmful to autonomy when someone *intends* to act against their best interest, although these instances of *akrasia* could be accounted for by deploying notice-as-friction.

Section 3.3 Designing a Value-Centered Privacy Assistant

So far, we have conceptualized what value-centered privacy choices are using 4DT and understood existing challenges – notice fatigue, lack of relevant controls, and (problematic) nudges – as a failure of the four-dimensions (Table 3-1). With this conceptual groundwork, we can now start *designing for* value-centered, autonomous decisions. In this section, we will utilize the dimensions of 4DT to systematically assess the degree to which current personalized privacy assistants (PPAs) create the space for privacy decisions that reflect our values (summarized in Table 3-2). In particular, we will assess the degree in which they address the challenges from the previous section and determine whether the design of PPA introduces any *new* autonomy frustrations that should be considered when designing a value-centered privacy assistant (VcPA) for smartphone app selection. We then propose three features for a VcPA based on this analysis – selective notices, exploratory notices, and suggesting alternatives. We then further visualize the VcPA features and the challenges of designing these features using user scenarios. In particular, we look at the tension between selective notices (*self-binding*) and exploratory notices (preventing the *inertia bias*), and how to tune the timing of exploratory notices in a manner that does not promote additional notice fatigue.

3.3.1 Evaluating Personalized Privacy Assistants (PPAs)

We can now turn to personalized privacy assistants (PPAs): what they are, and why explore them as a means of facilitating value-centered privacy choices in greater depth than explored in the Chapter 2. PPAs, currently under development by a team at Carnegie Mellon University,⁷⁸ would be machine learning assistants that personalize and automate privacy choices for a user. The team has explored PPAs to help manage user privacy in the Internet of Things (IoT) (Das et al., 2018) and smartphone applications (Liu et al., 2016), to name two. Personalized assistants could also utilize a variety of approaches to help a user manage their privacy, such as privacy-preserving nudges or semi-personalized setting recommendations based on user preference profiles (Liu et al., 2014; Story et al., 2020; Warberg et al., 2019). For example, Liu and colleagues (2016) designed a smartphone PPA

⁷⁸ Also see: *The Personalized Privacy Assistant Project* (<https://privacyassistant.org>).

where users got personalized recommendations for their privacy controls on their Android smartphone. They firstly developed privacy preference profiles based on a dataset of Android user privacy settings. During a user study of the PPA, participants took a dynamic, short quiz on their privacy preferences to sort them into a privacy preference profile. Based on their profile, they were then given recommendations on how they could change their privacy controls on their phone.

Conceptually, the development of PPAs has been fueled by a desire to help users make the best privacy choices for them in the current digital privacy environment. In particular, PPAs seem to be focused on overcoming one aspect of this larger issue – user notice fatigue (Liu et al., 2014). As Florian Schaub and colleagues have previously described (Schaub et al., 2015), determining the proper amount of notices is exceedingly difficult – too many notices causes users to simply “click through” them, and deploying them at too little a frequency does not provide the user with adequate information to make an informed decision. Through dynamic, personalized recommendation, PPAs aim to minimize notice by only presenting notices to the user that are relevant to them.

Besides minimizing notice fatigue, the personalization aspect of PPAs makes them an attractive technology to help us make more autonomous, value centered privacy decisions. As mentioned in Section 3.2, we need to be wary of deploying interventions – in this case, notices-as-friction – in an indiscriminate manner that may encourage value-inconsistent actions (decreasing *self-unification*). In addition, the personalization feature of PPAs makes them a better-suited technology for promoting value-centered privacy choices over privacy-enhancing technologies (PETs). As mentioned in Chapter 2, PETs aim to *protect privacy* through design choices rather than *promoting individual privacy decisions* that align with their privacy preferences (Garrido et al., 2022; Heurix et al., 2015). Similar to bright patterns, the privacy-by-design employed by PETs could result in a frustration of *self-unification* if one wishes to share data in a less privacy-restrictive way.⁷⁹

Using 4DT, we can now explore PPAs to assess to what extent they are successful at dealing with the issues of notice fatigue, the lack of relevant controls, and (problematic) nudges as outlined in Section 3.2. We can also use 4DT to identify additional areas of improvement to inform the design of a value-centered privacy assistant (VcPA), a system to help us make app choices in a manner that best reflects our personal values.

a 3.2.1 PPAs and Notice Fatigue

Not surprisingly, from the standpoint of 4DT, PPAs are mostly successful at addressing notice fatigue by utilizing *selective* (rather than general “blanket”) notifications (Liu et al., 2014, 2016; Warberg et al., 2019) (Table 3-2). As described in Section 3.2, the phenomenon of notice fatigue can be understood as a matter of degrees – a failure to *self-realize*, *self-unify*, *self-constitute*, or a combination of these. In summary, we can consider failures to: *self-realization*, where a someone “just clicks through” *all* notifications (with both *akratic* and *weakness of will* variations); *self-unification*, where this “click through” action does not match her personal identity and her values; and, in more extreme cases, *self-constitution*, where an “apathetic user” becomes so overwhelmed by privacy notices that they stop forming new intentions and commitments about her data privacy *at all*. In addition, challenges arise when two values come into conflict with each other. For

⁷⁹ There are, of course, strong normative arguments for privacy-by-design, and privacy-by-design was encoded into law by the GDPR (see: <https://gdpr-info.eu/issues/privacy-by-design/>). My aim here is not to discount the critical role privacy-by-design plays in broader, collective privacy considerations (especially those pertaining to data protection), but rather to identify an existing technology that could help promote individual value-centered privacy choices.

example, in the current privacy notice landscape, users who value *efficiency* could be in a *double bind* (fail to *self-unify*) if they also hold an opposing value in equal esteem.

Firstly, PPAs can act as a form of *self-binding* - that is, helping them follow through on their intentions. They do this by utilizing selective notifications to prompt the user to re-consider their privacy decisions that are at odds with their preferences. This, in turn, may help them act more consistently with their personal identity. By selectively “slowing down” the user in a manner tailored to their individual privacy preferences (Kahneman, 2011), they have further opportunity to pause and re-consider an intention that is at odds with the conclusion of her practical deliberation.⁸⁰ Friction in the case of PPAs can therefore be understood as a form of *self-binding* that helps the user *self-realize*⁸¹ and *self-unify* by triggering mindful reflection.

PPAs also introduce this friction in a manner that is more observant of what the user desires. I would argue that over-generous or designer-selective use of friction can cause problems, similar to dark patterns, when considered through a 4DT-lens.⁸² While these cases would encourage user reflection and, by extension, autonomous choice, the challenge here is that the *person designing the friction decides* when the user should slow down and be reflective, rather than when it would be most beneficial to the user to do so. This could either result in the user only making conscious decisions when the designer thinks that they should, or a well-intentioned designer may use friction too liberally as to cause notice fatigue. By selectively notifying users *based on what they think is best for them*, PPAs are theoretically better at assisting users without slipping into these issues.⁸³

Considering now the challenge of *double binds*, the PPA does help alleviate these conflicts between values – but only to a certain extent. In terms of the *double bind* concerning *efficiency* and other values described previously, the PPA could help the user focus their attention to decisions that are most relevant to their privacy preferences – thereby making the process more efficient. However, the number of selective notifications the user receives will still be inappropriately shaped by the *double bind* if the notices are deployed based on privacy preferences and not the values that define them. To demonstrate this, imagine the user who values *control* and *efficiency* equally engaging with a PPA system, which probes their privacy preferences and sorts them into a profile using a short privacy preference quiz. By answering these questions, the user is already forced to choose

⁸⁰ Recall from Chapter 2 that, by “slowing down,” I mean the shift from fast to slow thinking (Kahneman, 2011). “Fast” System 1 thinking is automatic, mindless, and ripe with heuristics and biases. “Slow” System 2 thinking is conscious, deliberate, and mindful.

⁸¹ In addition, Killmister (2017) notes that nudges cannot, generally, compensate for *weakness of will* (pg. 173), something I have mentioned as a problem when it comes to second-degree notice fatigue and being addressed here with selective notices. Friction is, of course, a nudge. However, when friction is used as a form of *self-binding*, it can also help us slow down and reconnect with our values – a “reminder” from ourselves of how to be and act – thereby helping us make more deliberate, conscious, value-centered choices. Therefore, I think selective friction can be used to promote value-centered data privacy decision making, even in *weakness of will* cases.

⁸² Of course, with better intentions. One such case is Terpstra et al. (2019), who I have been referencing here for their proposal to use friction to encourage reflection in privacy decisions. They are looking at using friction in a manner that encourages users to learn more about data privacy. While privacy literacy (its relevance and its promotion) is explored frequently in the privacy literature (e.g., Hagendorff (2018)), a 4DT lens suggests that such interventions could frustrate *self-realization* and *self-unification* (see discussion of nudges in Section 3.2). I am not intending to weigh in here on the merit of such an approach or on the greater privacy literacy debate, but rather, to explore whether it would be consistent with the value-centered, 4DT-informed understanding presented here.

⁸³ However, there are several concerns that PPAs need to consider as a machine learning assistant. These include minimizing bias and being sufficiently transparent. I discuss these issues in the next section when considering designing a PPA-like system for app selection, called a value-centered privacy assistant (or VcPA).

between *efficiency* and *control*. If they indicate that they want more restrictive privacy controls, they will receive more PPA notices than are consistent with their value of *efficiency*. If they answer the questions in a manner that prioritizes *efficiency*, they will not receive enough notices. In summary, PPAs help with *double binds* in part. However, basing notices on the values *behind the privacy preferences themselves* would likely help reduce *double binds* even more.

Table 3-2: PPA evaluation using 4DT and suggested modifications for a VcPA

Challenge	Relevant Dimensions	Do Current PPAs Address this Issue?	Suggested Modifications for a VcPA
Notice Fatigue	<i>Self-realization</i> <i>Self-unification</i> <i>Self-constitution</i>	In part: Personalized and selective notices help prevent <i>notice fatigue</i> and promote more value-consistent privacy choices through <i>self-binding</i> , but do not best prevent <i>double binds</i>	Keep selective notices, but base profiles on user values rather than privacy preferences alone to help minimize <i>double binds</i>
(Problematic) Nudges	<i>Self-realization</i> <i>Self-unification</i>	No: the <i>inertia bias</i> makes it difficult for users to change from their initial privacy settings	An <i>exploratory process</i> using <i>exploratory notices</i> , that does not encourage one download choice over another
Lack of Controls	<i>Self-realization</i> <i>Self-unification</i>	No: because of the <i>limited granularity problem</i> , a user's ability to realize their commitments and act on them is hindered	<i>Suggesting alternatives</i> to quickly link users to similar applications whose data collection practices better align with their values

b 3.3.2 – Other Challenges: Inertia Bias and Privacy Controls

Like the issue of notice fatigue is *inertia bias* - a kind of nudge, in this case. As mentioned in the Chapter 2, this cognitive bias makes it difficult for users to change from their initial privacy settings – even if they wish to (Thaler & Sunstein, 2008). In the case of PPAs, this challenge continues to manifest as a failure to update and change their privacy profile, even if they are given the ability to do so. This could constitute an autonomy violating “nudge” under 4DT. Returning to the Liu and colleagues (2016) work on PPAs, these researchers used nudges to test if users would change their profile. Most did not. In addition, while users were always able to change their profile, few did. While the authors claimed that this supports the accuracy of their profiles, this could also be interpreted as evidence of the *inertia bias*. We can also imagine someone who continued to receive notifications based on their original privacy profile and acted according to them even if it was determined that this profile is not the best fit for them. Like certain manifestations of notice fatigue, failing to act on this intention would result in a failure to externally *self-realize* (if *deliberative*). In addition, their resulting privacy decisions would likely not be consistent with their values, a violation of *self-unification* (both *deliberative* and *non-deliberative*). If PPAs do not account for the *inertia bias*, they cannot be said to be fully *self-realizing* or *self-unifying*.

Lastly, there is the challenge at the level of privacy controls themselves. As described in Section 3.2, there are many different privacy notice designs with different degrees of privacy control (Utz et al., 2019). Smartphone PPAs can also only modify preferences using the smartphone operating system's (OS) available controls. These controls (e.g., access to Contacts, Camera, Location) have been previously shown to be insufficient for capturing the privacy concerns of users (Felt et al., 2012). Terpstra and colleagues (2019) have proposed that a lack of meaningful privacy controls in conjunction with positive friction – in this case, a notice from the PPA – can make us frustrated, possibly undermining positive friction benefits such as value reflection. This frustration can be understood in terms of failures to *self-realization* and *self-unification*. If a user forms an intention (privacy preference) but the control is not present, this would be a frustration of *self-realization*. Because they are not acting according to their values, this would also constitute a failure of *self-unification*. Due to the limited level of granularity available to them, their ability to realize their commitments by making consistent choices is still limited when using a PPA.

3.3.2 Considerations for Selecting Smartphone Applications

These challenges present three major implications for designing a PPA-like system, a value-centered privacy assistant (VcPA), to help users select smartphone applications. Firstly, like current PPAs, users of a VcPA may fail to update and change their privacy profile due to the *inertia bias*, hindering *self-realization* and *self-unification*. Secondly, these dimensions could be further undermined due to the lack of granularity problem. While the challenges to *self-realization* and *self-unification* of current PPAs are a result of the lack of granularity of either 1.) privacy notice design (online PPAs) or 2.) the OS' privacy setting controls, a PPA-like system assisting with smartphone application selection is essentially a dichotomous decision: either download, or not download. Not only does this minimize *self-realization*, but this will prevent a user from realizing their commitments in a manner consistent (*self-unifying*) with their values by introducing additional or aggravating current *double bind* situations (such as the choice between two social media apps that may not have the same *social connection* value). Thirdly, the aforementioned *double bind* situation concerning *efficiency* and PPAs will also need to be considered when selecting an app; someone who values *efficiency* may still face instances of *double binds* if selective notifications are based on the app's data collection practices and her privacy preferences alone.

3.3.3 Suggested Modifications when Developing a Value-Centered Smartphone Privacy Assistant

There are a few possible modifications to current PPA design that could overcome these remaining challenges and create a value-centered privacy assistant (VcPA) (Table 3-2).

To rectify the *efficiency* value challenge, VcPA profiles could instead be based on the user's personal values. By basing profiles on values rather than privacy preferences alone, the VcPA could be more accurately tuned to prevent *efficiency*-based *double bind* situations and maximize *efficiency*. Critically, however, user tests are required to: 1.) determine in what way values intersect with app data collection preferences; and 2.) how this intersection could be operationalized as VcPA profiles. In addition, the profiles must also: 1.) be an *accurate* reflection of a user's values; 2.) have an accessible and *understandable description*; 3.) clearly state *how* the privacy preferences and values are utilized in profile creation and assignment; and 4.) be able to be *changed* to a different

profile if notices are unhelpful. All these considerations must be accounted for when designing and testing a VcPA prototype to rectify the *double bind* challenge without sacrificing accuracy and transparency. In addition, explaining why they are receiving the notice will be critical to deal with users who may be exhibiting *akrasia*, where they have made a deliberate choice to act against their values and where the lack of such an explanation could result in a nudge that frustrates *self-realization*.⁸⁴

Self-realization and *self-unification* as defined by 4DT are minimized in part because PPAs are limited by the lack of relevant controls, which remain relevant to VcPAs, through a download-or-not-choice that may make acting consistently with their values in such a binary situation not possible. This challenge could be overcome in a VcPA by not only informing the user when an application is requesting privacy settings that are *inconsistent* with their values, but by also suggesting alternative applications to quickly link them out to similar applications that better align with their commitments and values. This increases the likelihood that they will be able to find an app whose data collection practices match their values, thus (better) upholding *self-realization* and *self-unification*. In the best-case scenario, such an app would also not cause the user to compromise one of their values in the case of *double binds*.⁸⁵

To be more fully *self-realizing* and *self-unifying*, however, a VcPA will also have to tackle the *inertia bias*. Agreeing with others that there is a need for “learning” (or, at least room for change) in the privacy notice process (Terpstra et al., 2019) this could be tackled by periodically (but not excessively) “mining” user goals, values, and preferences – that is, checking in with the user that the notices are still a good fit for them by making them aware, perhaps by using an *exploratory notice*. It will be critical, however, that this “mining” process be sufficiently random and dispersed as to *not* encourage user action one way or another. If a user makes a choice that is inconsistent with what they have resolved to do because of this exploratory notice, this would be inconsistent with *external self-realization* (and likely *self-unification* if the action goes against their values). This will be particularly applicable to those who are, for whatever reason, authentically *not as concerned about privacy* (that is, most data sharing is aligned with their values).⁸⁶ It is, fundamentally, a tension between *self-binding and exploration*; nudging in a value-aligned manner, or against. The key to managing this will likely be the same as dealing with *akrasia*. It will be critical to be fully transparent in our explanation of why the user is receiving the exploratory notice – that is, to encourage System 2 (deep, deliberative) thinking through friction rather than “just nudging” (Kahneman, 2011).

a Further Visualizing a VcPA Using (Tentative) User Groups

To further demonstrate these challenges as well as VcPA features, let us consider a few scenarios. We can consider high-level user scenarios using Westin’s privacy groups (Hoofnagle & Urban, 2014).⁸⁷ These groups are: privacy fundamentalists, or users who are very concerned about disclosing their data even in the presence of privacy protections; privacy pragmatists, or users who have very specific privacy concerns about data disclosure in certain contexts; and the privacy unconcerned, or users who have mild or no

⁸⁴ See the discussion of nudges in Section 3.2 and footnote 71.

⁸⁵ See example of Shauna in Table 3-1, who values both *social connection* and *trustworthiness* when engaging with a messaging app.

⁸⁶ See the example of Ash in Table 3-1, who believes in sharing their data with companies to promote *innovation*.

⁸⁷ Which are, as explored in the introduction (Section 2.2), highly flawed. Here, they are used as a rough, preliminary tool to further visualize the VcPA and the challenges of making one in specific scenarios.

concern about disclosing data. User scenarios are common in user-centered design, where designers can utilize scenarios as a means of translating high-level ideas into more concrete possibilities. For our purposes, I define user scenarios as “narrative descriptions” of a user’s engagement with a VcPA (Rosson & Carroll, 2002). In particular, these user scenarios have a descriptive emphasis on the user’s goals and values to demonstrate the key features and challenges of making a VcPA.

In each scenario, all three hypothetical users are faced with the decision whether to download the application OpenLitterMap.⁸⁸ OpenLitterMap is a citizen science initiative that allows users to take smartphone pictures of litter and upload them into a publicly available dataset. The goal is to empower citizens to be active participants in combating local pollution. Photos of litter can be uploaded anonymously or with a username to participate in the litter “World Cup.” In both cases, the system records several features, including time, date, location, and phone model. This could have many implications, of which I will focus on one. This data sharing means that in areas of low app use, it becomes possible to identify a user based on inference. From a value-centered privacy approach, a potential OpenLitterMap user will need to balance the value of disclosing information against the possible (albeit, small) risk of identification.

(1) User #1: Privacy Fundamentalist

Firstly, let us consider User #1 – the privacy fundamentalist. User #1 likes to make environmentally friendly choices. They are willing to do what they can to preserve the environment and provide the best future for their children. User #1 hears about OpenLitterMap from a friend and goes to download it, acting quickly without deliberating. With the VcPA system, a notice appears on their screen, warning them that this application is not consistent with their personal values of *security* and *control*. They decide to check out other apps first by clicking “see alternative applications.”

At this point, there are two possible outcomes for User #1. The first is that they find a different litter clean-up application that is consistent with their values of *security* and *control* and download that one instead. In this application, the data collected may, for example, only be accessible to policy makers and environmental scientists, be encrypted, and not collect their phone model. The second possible outcome is that User #1 does not find another application with a similar function. They may then decide to stick to their regular beach cleanings to help their environment instead of downloading an application.

Without the VcPA system, User #1 may click through the privacy settings and allow the app to access their photos, camera, location, date, time, and phone model. They begin using the app when they are walking to pick up their children from school. While they want to help document litter and believe in allowing data scientists access to their documented litter data for environmental research purposes, they would feel uncomfortable if someone was able to identify their route to and from the school – and, by association, information about their children. To uphold their values of *security* and *control*, they may decide to upload their litter anonymously rather than with a username. However, they may be the only one using OpenLitterMap on that route, and it would be possible for someone looking at the data to identify them. While some may have been comfortable with this level of risk, they would not have been – they prioritize *security* and *control* over their value of *environmentalism*.

In this first user scenario, the absence of the VcPA would have resulted in a violation of their *self-unification* – their actions (to download the application) would not be

⁸⁸ <https://openlittermap.com>

in alignment with their values (*security* and *control*). Thanks to selective notifications, however, User #1 is alerted to this misalignment of their action and their values. In addition, their *self-realization* (acting on their values) could also be enhanced if they are able to find another app using the “suggest alternatives” feature following further deliberation.

There are, however, two important caveats to also consider. If User #1 is, instead, *akratic* – that is, deliberated and formed the intention not to act according to their values – such a notice would slightly diminish their *self-realization*. That said, as stated in 2.2.2 and 2.3.3, we could introduce notices with an explanation of why they are receiving the notice. Then, we can only hope that our reasoning causes User #1 to reconsider and form new intentions. It is also possible that User #1’s values have changed, and they are no longer receiving notices in line with their current value set. In this case, the selective notice may encourage them to take an action that is not *self-unifying* – that is, consistent with their values. This is an example of the tricky balance between *self-binding* and the *inertia bias*, which must be accounted for in exploratory notice timing.

(2) User #2: Privacy Pragmatist

VcPA Users #2 and #3 are quite similar when viewed through a 4DT value-centered lens – demonstrating the tricky balance between accounting for the *inertia bias* and inappropriate nudging, as well as establishing optimal exploratory notice timing. Let us start with User #2. This user, the privacy pragmatist, has a few practical apps on their phone. A colleague recommends that they take a look at OpenLitterMap. They go to download it.

With the VcPA, there would be two possible outcomes for User #2. They could firstly receive an exploratory notice letting them know that, while this application is consistent with their previously stated values, there is a chance of violating the values of *security* and *control* if they use the application. User #2 will then have to decide whether to download this application when faced with this new information. In the absence of an exploratory notice, User #2 may simply click through the privacy settings and allow the app to access their photos, camera, location, date, time, and phone model, the same result without the VcPA system.

In this case, User #2’s autonomy may be enhanced with the VcPA. The exploratory notice could help fight the *inertia bias* if their *self-defining* attitudes (values, goals, beliefs, commitments) have shifted since choosing their initial VcPA profile. If, however, they have not changed their values and they do not download it because of the added notice, this would actually hinder *self-unification* because they would act in a manner inconsistent with their values. In addition, if User #2 has *deliberated* and decided to download OpenLitterMap and does not download it, we have also frustrated their *self-realization*. In this manner, introducing added friction in the form of an added notice, *self-realization* is reduced by providing a barrier to realizing their values and intention. However, *not* having an exploratory notice means allowing the potential *inertia bias* to go unchecked. If User #2 receives an exploratory notice letting them know that, while this application is consistent with their previously stated values, there is a chance of violating their values if they decide to alter their actions, we may encourage them to act against their values if their values have not changed. In this case, we would have caused a failure of *self-unification* and possibly *self-realization*. As we will see, this failure could also occur with User #3, and accounting for it will be critical for both hypothetical users.

(3) *User #3: The Privacy Unconcerned*

Lastly, let us consider User #3, the privacy unconcerned.⁸⁹ User #3 attends a talk organized by their local Greens Club about the harmful effects of litter. The Greens Club recommends checking out OpenLitterMap. User #3 likes the idea of creating a profile to compete for the OpenLitterMap “World Cup” leaderboards. They go to download the application.

If they do not receive an exploratory notice, it is likely that they will download the application anyway, regardless of whether their values have changed (due to the *inertia bias*). Like User #2, if they receive an exploratory notice,⁹⁰ they decide to alter their actions, and their values have *not* changed, we have caused a failure of *self-unification* and possibly *self-realization*. However, in the absence of an exploratory notice, the *inertia bias* could still be an issue.

In both of these cases, I would argue that the key to resolving this tension is in the explanation on the exploratory notice and our desire to trigger reflection through friction, rather than “just nudging.” Again, like User #2, we can hope, based on the text of the exploratory notice (“while this application is consistent with your previously stated values, there is a chance of violating the values of X and Y if you use this application”) that the rationale will prevent them from not downloading the app if it is fully consistent with their values to do so. This better upholds *self-realization* and *self-unification* while allowing us to also account for the *inertia bias*.

However, finding the optimal timing for the exploratory notice will also be critical, as too much “exploration” with pop-up notice could lead to (third-degree) notice fatigue. This would, essentially, bring us back to square one – too many notices creating “apathetic users” who decide to no longer take on commitments pertaining to their data privacy.

(4) *Conclusions from User Scenarios: A Difficult Balance*

As these examples demonstrate, the execution of the notices will be particularly difficult depending on whether the user’s values have shifted and what action the notice ultimately causes them to take. While the VcPA can act as a form of *self-binding* and promoting value-centered privacy choices (e.g., for User #1), it could encourage an action inconsistent with their values if their values have changed. To counter this, we have exploratory notices, where we have an explanation which, like the explanation on the selective notice to “reason” with the *akratic*, hopefully prevents us from unintentionally nudging someone to take a *non-self-unifying* action. However, the timing of this notice will be difficult to tune, as too much “exploration” could lead to notice fatigue. Striking this balance will require “tuning” the timing of the exploratory notice’s based on empirical data (user studies).

Section 3.4 Conclusion

Here, I have explored how we can *create the space* for value-centered privacy decisions by applying 4DT. I first conceptualized privacy decisions in terms of the four dimensions –

⁸⁹ For the sake of argument, I am assuming that User #3 is “genuinely” unconcerned about privacy – that is, not “apathetic,” but that sharing data is *truly* consistent with their values. See Section 3.2.2., footnote 64, or the example of Ash in Table 3-1 (an agent who values *innovation* and believes in sharing their data with companies to improve the quality of services provided).

⁹⁰ E.g., an exploratory notice letting them know that, while this application is consistent with their previously stated values, they would be violating the values of *security* and *control* if they decide to download the app.

self-definition, self-realization, self-unification, and self-constitution – and explored existing data privacy challenges through this lens. To inform the design of a smartphone assistant that creates this space for users, I next examined PPA technology using a 4DT lens. While PPAs, with their use of selective notices, are partly successful at creating the space for value-centered choice, several concerns around the *inertia bias, double binds,* and lack of controls remain. Using these insights, I lastly propose a value-centered, smartphone privacy assistant, (VcPA) to help users make more value-centered decisions at one privacy decision point: smartphone app choices. To best promote value-centered choice, a VcPA can keep the selective notices of the original PPA (based on values *and* privacy preferences) and add “suggest alternatives” and “exploratory notice” features. While the tension between selective notices (*self-binding*) and exploratory notices (preventing *inertia bias*) will be difficult to rectify, we can aim to do so by making sure our exploratory notice application is appropriately explanatory. However, we will need to be careful tuning the timing of these notices appropriately to not lead to (more) notice fatigue.

Empirical studies, however, are required to validate this conceptualization of value-centered data privacy decisions and to assess whether a VcPA is effective at promoting value-centered choices. In the next three chapters, I describe a mixed-methods study that aims to do just that.

Chapter 4 Methodology for Empirical Studies

*Code Monkey get up, get coffee
Code Monkey go to job
Code Monkey have boring meeting
With boring manager Rob
Rob say Code Monkey very diligent
But his output stink
His code not "functional" or "elegant"
What do Code Monkey think?*

Jonathan Coulton (“Code Monkey”)

Section 4.1 Chapter Overview

In this chapter, I present the mixed methods, empirical study design utilized to empirically assess the role of values in privacy decisions in order to promote value-centered choices and establish the usability and effectiveness of the value-centered privacy approach described in Chapter 3. To accomplish this, we aimed to 1.) empirically explore the relationship between a user’s personal values, privacy preferences, and smartphone app choices; and 2.) design and test a prototype value-centered privacy assistant (VcPA).

The study consisted of three phases (Table 4-1). Phase I involved an online survey of values, privacy preferences, and smartphone apps. This provided quantitative data scoring values as overall life-guiding principles; scoring values when deciding whether to download a specific app; and binary indicator of privacy preferences.

Phase II involved testing a prototype VcPA system informed by Phase I results. To accomplish this, a testing environment – called the Mock App Store (MAS) – was designed for testing the VcPA. The MAS is a web interface that replicates certain features of the Apple App Store and includes a “virtual” smartphone to “download” apps. Participants in Phase II were asked to browse the MAS and download apps. The system recorded interactions with the VcPA and the Store, as well as elicited feedback on VcPA features (e.g., selective notices, exploratory notices, and a “suggest alternatives” page; see Chapter 3, Table 3-2).

To provide further depth in our exploration of the value-privacy relationship, Phase III consisted of follow-up semi-structured interviews with some Phase II participants. These interviews probed store participants values, privacy preferences, and app choices on the MAS as well as in their everyday life.

The three phases were integrated using a process of convergent design. Convergent design involves bringing together results obtained via different methods to answer a common question (Fetters et al., 2013). Two research questions were posed to guide design and analysis:

1. **RQ1:** *What is the relationship between values and privacy preferences when deciding to download an app, if any?*

2. **RQ2:** *How useful and effective is a value-centered privacy assistant at helping users make app choices consistent with their values?*

In this case, the online privacy preference and value survey (Phase I) were primarily aimed at answering **RQ1** and the VcPA user study (Phase II) at **RQ2**, with the interviews (Phase III) containing questions pertaining to both research questions (Table 4-1). Phase I was also completed first and used to inform the design of Phase II and Phase III. Initial survey results informed interview question selection, and the value-privacy profiles for the prototype VcPA were derived from the survey data.

Table 4-1: Mixed-methods empirical study design, in three phases

	<i>Brief Description</i>	<i>Timeline</i>	<i>RQ1: What is the relationship between values, privacy preferences/choices, and app selection, if any?</i>	<i>RQ2: How well does a value-centered privacy assistant help users make choices more consistent with their values?</i>
<i>Phase I. Online Value and Privacy Preference Survey</i>	A survey of user values and smartphone app privacy preferences	October – December 2021	Are there any observable and/or statistically significant correlations between different apps, privacy preferences, and/or values?	
<i>Phase II. Mock App Store Study</i>	A user test of the VcPA in a synthetic, online smartphone app store	June 2022- October 2022		How did participants interact with the VcPA? What could be improved?
<i>Phase III. Post-Study Interviews</i>	Semi-structured interviews with some Mock App Store participants exploring how they interacted in the VcPA and navigate data privacy choices “in real life”	July 2022- October 2022	How do users care about data privacy and smartphone app choice? What values are involved and in tension?	How did participants find the VcPA prototype? What could be improved?

4.1.1 Collaborator Contributions

The empirical study design presented in this chapter was done in collaboration with Prof. Dr. Mathieu d’Aquin (supervision guidance, study design feedback, data analysis, Mock App Store implementation), Dr. Heike Felzmann (supervision guidance, study design feedback, interview analysis, research ethics application feedback), Prof. Dr. Kathryn Cormican (supervision guidance, study design feedback), Dr. Dave Lewis (supervision guidance, study design feedback), Dr. Ilaria Tiddi (supervision guidance, Mock App Store implementation), and Dr. Dayana Spagnuolo (data analysis, value-centered privacy assistant implementation). I (the PhD candidate) primarily designed and conducted the study, working with collaborators at all stages of design, data collection, and data analysis.

4.1.2 Relevant Papers and Conference Contributions

Some material in this chapter, including certain text and figures, has been previously published or presented in the following:

Carter, Sarah E., & Felzmann, Heike. (2023). How do we value data privacy? Insights and design implications. To be published in: *Engineering and Value Change* (part of: *Springer Philosophy of Engineering and Technology series*). Abstract available at: <https://zenodo.org/record/8367542>

Carter, S.E., d’Aquin, M., Spagnuolo, D., Tiddi, I., Felzmann, H., & Cormican K. (2023). The privacy-value-app relationship and the value-centered privacy assistant. ArXiv. <https://arxiv.org/abs/2308.05700>

Carter, Sarah E., & Felzmann, Heike. (2023, April 21). How do we value data privacy? Initial results from semi-structured interviews. Forum on Philosophy, Engineering, and Technology (fPET2023), Delft, the Netherlands. Zenodo. <https://doi.org/10.5281/zenodo.8204406>

Carter, Sarah E., Tiddi, Ilaria, & Spagnuolo, Dayana. (2022, June 13). A "Mock App Store" interface for virtual privacy assistants. Hybrid Human Intelligence 2022: Augmenting Human Intellect (HHAI2022), Amsterdam, the Netherlands. Zenodo. <https://doi.org/10.5281/zenodo.8204393>

Carter, S. E., Tiddi, I., & Spagnuolo, D. (2022). A “Mock App Store” interface for virtual privacy assistants. In S. Schlobach, M. Pérez-Ortiz, & M. Tielman (Eds.), *HHAI2022: Augmenting Human Intellect* (Vol. 354). IOS Press. <https://doi.org/10.3233/FAIA220212>

Carter, Sarah E. (2021, September 1). PhD proposal: conceptualizing and realizing a value-centered privacy assistant. Doctoral Symposium: The 20th IFIP Conference e-Business, e-Services, and e-Society (I3E2021), Galway, Ireland. Zenodo. <https://doi.org/10.5281/zenodo.8204916>

4.1.3 Research Ethics Approval

The study and any amendments were approved by the National University of Ireland – Galway (now University of Galway) Ethical Review Committee prior to participant recruitment and in accordance with all university policies.⁹¹

Section 4.2 Phase I: Online Value and Privacy Preference Survey

4.2.1 Phase I Summary

The online value and privacy preference survey aimed to evaluate the relationship between user values, smartphone app choice, and privacy preferences (**RQ1**; Table 4-1). In particular, the survey aimed to explore whether there are any *quantitatively observable or*

⁹¹ <https://www.universityofgalway.ie/research-office/policiesandprocedures/>

statistically significant relationships between basic human motivational values, smartphone privacy preferences, and the decision whether to download an app. Understanding how values overlap with certain app permissions in different contexts was also utilized to develop profiles for the VcPA in the Mock App Store Study (Phase II). The survey was conducted from October-December 2021.

4.2.2 Online Survey Structure

a Survey Design Overview

The online survey began with demographic questions (age, gender, nationality, English proficiency, smartphone use, and education). It then asked participants to rank the importance of ten values on a scale of 1 to 9 in their lives more broadly and in the specific context of selecting an app. Participants were also asked about their app privacy preferences. More details on how values and privacy preferences were scored are provided in (1) and (2), respectively, along with other relevant survey details in (3-5).

(1) Measuring Values: Applying the Short Schwartz Value Survey (SSVS)

Values in the survey were measured using the Short Schwartz Value Survey (SSVS) (Lindeman & Verkasalo, 2010), slightly modified for our purposes. This survey asks participants to rank the importance of ten values on a scale of 0 (Opposed) to 8 (Of supreme importance). This survey is theoretically grounded in the Theory of Basic Human Values (TBHV),⁹² where values are understood as motivators that drive us to meet certain human or societal needs (Schwartz, 1992, 2012). Because the original 56-value item SVS questionnaire⁹³ utilized in the TBHV can be time-consuming and impractical for certain situations, Lindeman and Verkasalo (2010) designed the Short Schwartz Values Survey (SSVS). The SSVS consists of only 10 questions asking participants directly to rank their 10 value items on a scale of -1 (against my principles) to 5 (of supreme importance). For our purposes, the SSVS was selected because we required a validated method of measuring values in survey form that was sufficiently brief as to not overwhelm participants. We slightly modified the SSVS by changing the scale to 1 (opposed) to 9 (of supreme importance based on survey pre-testing feedback (see (4))). We also asked participants to rank their values twice: once as broad, life-guiding principles as in the original SVS/SSVS, and once in the specific context of selecting an app. This resulted in 20 value questions.⁹⁴

(2) Evaluating Privacy Preferences: Utilizing Apple Privacy Labels

Survey questions pertaining to smartphone data collection practices were based on Apple's App Privacy Labels,⁹⁵ which categorize application data collection practices according to

⁹² The theoretical background of the TBHV, as well as more details on the SVS, are presented in Chapter 2.

⁹³ An updated SVS is also available that explores 19 broad values based on sub-values observed within in the original ten. For example, Schwartz and colleagues (2012) proposed further subdividing security into societal security and personal security based on emerging empirical evidence. However, the 19 values are compatible with the 10 original motivational theories, as they are a finer-grained version of the original ten.

⁹⁴ Survey questions available in Appendix I.

⁹⁵ <https://www.apple.com/privacy/labels/>

type and linkage to a user. This resulted in three questions for each data type (e.g., *location*) about a participant's permission comfort level of each type of data when it is 1.) linked to them, 2.) not linked to them, and 3.) used for tracking (Table 4-2) (e.g., *tracked location*). Apple privacy permissions were chosen over Android because of feasibility, again desiring to focus on what preferences are likely to be the most relevant and not overwhelm participants with a long survey.⁹⁶

(3) *Apps Explored*

Previous studies on both personalized privacy assistants (PPAs) and app design suggest that both values and privacy preferences vary depending on the app context (Liu et al., 2016; Nurwidiantoro et al., 2022). To determine whether values and permissions varied based on context in our study, two versions of the surveys were created in one of two contexts: 1.) health and fitness apps (Lose It!) and 2.) environmental apps (OpenLitterMap).⁹⁷ In greater detail, Lose It! is a health and wellness app that helps you track your eating habits and exercise to meet weight loss goals. OpenLitterMap is a citizen science app that allows a user to upload pictures of litter in their community to create a publicly available dataset. A participant was randomly assigned a version of the survey.⁹⁸

(4) *Survey Feedback and Pre-Testing*

Survey feedback and testing aimed to encompass many disciplines due to the interdisciplinary nature of the study. Feedback on the survey design was given in an iterative manner by PhD supervisors Prof. Dr. Kathryn Cormican, Prof. Dr. Mathieu d'Aquin, Dr. Heike Felzmann, and Dr. Dave Lewis. Dr. Clare O'Dwyer, a postdoctoral researcher under Prof. Dr. Cormican, also provided significant feedback on statistical considerations. Graduate Research Committee (GRC) members (Prof. Dr. John Breslin, Dr. Karen Young, and Dr. John Danaher) provided additional feedback during their yearly PhD project review. The survey was pre-tested with a group of multi-disciplinary colleagues from the National University of Ireland – Galway (now University of Galway). Cumulatively, these individuals had a broad range of expertise, encompassing data science, law, engineering, human-computer interaction (HCI), quantitative and qualitative methodology, applied ethics, and philosophy.

Based on survey pre-testing and expert feedback, the survey scale was modified from the original SVSS scale of -1 - 5 to a scale of 1 - 9 to add an intuitive midpoint at the number "5." In addition, there were a few small, mostly cosmetic, alterations to the survey to make it more attractive and understandable for participants.

⁹⁶ There are 40 different Apple privacy permissions based on Apple Privacy Labels (Table 4-2). At the time this study was designed, Android did not have a simplified, accessible list of app privacy permissions, and there were over 300+ Android privacy permission without a simplified privacy label ontology (permissions viewable at: <https://gist.github.com/Arinerron/1bcaadc7b1cbeae77de0263f4e15156f>). The Google Play Store released their own privacy labels (remarkably similar to Apple's) just as this study was coming to a close (Velazco, 2022).

⁹⁷ Lose It!: <https://www.loseit.com/>; OpenLitterMap: see footnote 88.

⁹⁸ Survey copies available in Appendix I.

Table 4-2: Apple Privacy Label options⁹⁹

Linkage	Unlinked	Linked	Tracked
Data Type	<ul style="list-style-type: none"> • Your health and fitness information (such as your health data and exercise data) • Your financial information (such as your credit card number, bank account number, form of payment, credit score, salary, income, and debts) • Your location • Your sensitive information (such as your race, ethnicity, religion, political affiliation, or sexual orientation) • Your contacts (such as your address book) • The content of your phone (such as emails, text message, photos, and audio data) • Your browsing history (such as when you are browsing a website, outside of the app) • Your search history in the app • Your purchase history in the app • Your usage data (such as clicks, scrolls, taps, or advertisement views) • General diagnostic data (such as crash logs, launch time, and energy use) • None 		
		<ul style="list-style-type: none"> • Your contact information (such as your name, email address, phone number, or physical address) • Other identifiers (such as your account username or the identifier for your phone) 	

4.2.3 Participant Recruitment

To participate in the survey, participants needed to be: 1.) aged 18 years or older (for consent purposes); 2.) a native or fluent speaker of English (for survey understandability); and 3.) had owned or currently owned a smartphone (for study relevance). Based on a simple sample size calculation with a 95% confidence interval and a 5% margin of error for the very large population of interest (English speaking smartphone users over the age of 18), we aimed to recruit ~300 participants. These 300 would be split evenly between the Lose It! and OpenLitterMap versions of the survey. The survey was implemented using Microsoft Forms and responses were collected using snowball sampling. The survey was distributed via social media (Twitter, LinkedIn, and Facebook) and the researchers’ professional networks (consisting primarily of academics in Ireland, the United States, and the Netherlands) from October-December 2021 (Table 4-1). Consent and the study information sheet were made available when first opening the Microsoft Form. Responses from the survey were anonymous, linked only to a numeric identifier.

⁹⁹ See footnote 96 for more on Apple Privacy Labels.

4.2.4 Statistical Analysis: Value-Privacy Preference Correlation

Data was first cleaned to create consistency between privacy preferences within each data type based on the most privacy invasive preferences. For example, if participants stated that they were comfortable with *tracked location* data being collected, we checked that they also found *linked location* and *unlinked location* data acceptable.¹⁰⁰ Correlations between permissions, the app being considered, and value scores were calculated using Spearman correlations. In particular, Spearman's rank correlation coefficients (ρ) and the significance of each correlation (p) was calculated between app-specific values scores (*Value App*), general life-guiding value scores (*Value*), the difference between *Value* and *Value App* (*Value Difference*), the app being considered (*LoseIt!* or *OpenLitterMap*), and privacy preference (e.g. *tracked location*).

Section 4.3 Phase 2: Mock App Store Study

4.3.1 Phase II Summary

Following the online survey, a second, smaller group of users were recruited to partake in a proof-of-concept study exploring the usability of a value-centered privacy assistant (VcPA) (RQ2; Table 4-1). This study, called the Mock App Store Study, aimed to assess whether a prototype VcPA system increases app choices more consistent with a user's values and whether users find the system helpful. These users engaged with the VcPA in an online simulation of an app store, the "Mock App Store" (MAS) environment. Participants were also asked to complete a survey before and after the exercise on their privacy attitudes and experience with the VcPA.

4.3.2 Mock App Store and VcPA Design

The Mock App Store (MAS) consisted of an interface designed to look like and emulate the function of an app store, including "download" apps; and the VcPA system (consisting of selective notices, exploratory notices, and a "suggest alternatives" feature – see Table 3-2 for VcPA features). Specifics about the MAS and VcPA design are provided in the following subsections (a-e), and a demonstration of the MAS with basic VcPA features can be viewed at: <https://youtu.be/ziGoowteN6E>.

a Mock App Store (MAS) Development

The MAS includes 97 apps from the health and fitness category extracted from the Apple App Store (Figure 4-1). The 97 apps were derived by first selecting the top ten apps from the US App store on Sept. 29th, 2021. The app IDs for these apps were then used to submit an API request for the top 10 similar apps¹⁰¹ to AppTweak.¹⁰² The starting app and its

¹⁰⁰ We also considered the correction in the opposite direction (e.g., no *unlinked location* → no *tracked location*) but decided against it. How the survey was designed (asking participants to check a box next to data they would be comfortable sharing) makes the "default" setting for each privacy preference a "no" answer. Participants had to explicitly select and check boxes. In instances of inconsistency, such as a checked *tracked location* and unchecked *unlinked location* and *linked location*, it is therefore more likely that they missed checking all three levels (unlinked, linked, and tracking) for one data type than accidentally checking one.

¹⁰¹ Command: *similar apps, outgoing*. This shows the apps that currently appear directly on the app's similar section on the US Apple App Store.

¹⁰² <https://www.apptweak.io>

similar apps were then grouped together to make an “app family.” This resulted in 110 apps sorted into 10 families of 11 apps each. Then, the metadata¹⁰³ for all apps was requested from AppTweak. The top 15 keywords for the apps were also requested from AppTweak and added to the metadata .csv file. Data was then cleaned to remove non-alphanumerical characters and for ease of processing.

To deal with duplicate apps in different families, we firstly merged app families that contained many apps in common. To determine the best fit for the other duplicates, we used the Jaccard index based on app keywords (Jaccard, 1901; Niwattanakul et al., 2013). Finally, apps that were geography, occupation, or product dependent were removed. To increase app diversity and to replace the apps that were removed, we repeated the process for 3 additional apps: Down Dog (a yoga app); Headspace (a popular meditation app); and Pray.com (Christian prayer and Bible study app). This resulted in a total of 97 apps, sorted into 9 families (Appendix II). This data was then utilized to generate apps on the MAS interface.

The logic for the Mock App Store was coded using Python Flask¹⁰⁴ and integrated into the app store interface using HTML. The MAS itself was designed in Javascript.¹⁰⁵

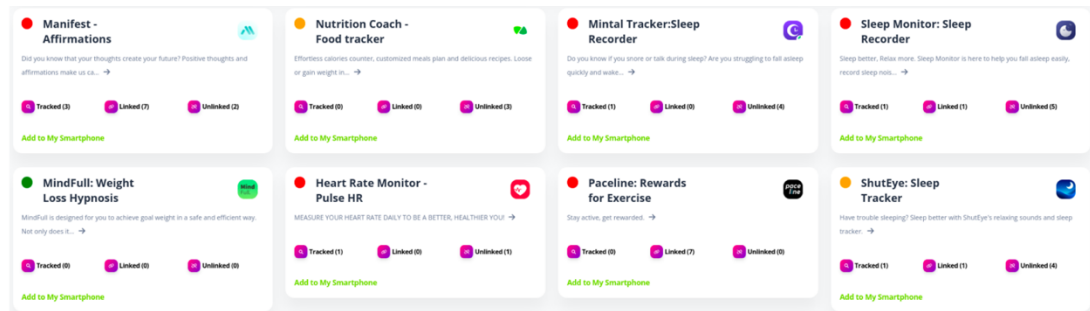


Figure 4-1: Mock App Store main page with a few example apps

b VcPA Profiles

Profiles for the VcPA were designed based on the clusters identified by hierarchical clustering. Hierarchical clustering was conducted on app-specific values and stated general life-guiding principles.¹⁰⁶ Hierarchical clustering has been used previously to identify privacy profiles in smartphone privacy assistant settings (Liu et al., 2016). The app-specific clustering yielded two clusters of those who care about every value very highly and those who do not care much about any value. Hierarchical clustering based on general values proved more promising. Following z-score normalization,¹⁰⁷ hierarchical clustering yielded three clusters.¹⁰⁸ Clustering was then verified by graphing profiles in three dimensions according to the variables (*Power*, *Hedonism*, and *Achievement*) with the highest variance (Figure 4-2). Visual analysis of the corresponding 3D plot suggested that cluster 1 (red)

¹⁰³ The metadata includes the app’s description, title, icon, and app ID as listed in the last 30 days on the US Apple App Store.

¹⁰⁴ Python Flask: <https://flask.palletsprojects.com/en/2.3.x/>

¹⁰⁵ Video demonstration of the MAS: <https://youtu.be/ziGoowteN6E>

¹⁰⁶ K-means clustering was also considered but was not selected because it required that we set the number of clusters.

¹⁰⁷ Standardization such as z-score normalization is used in cluster analysis because clustering algorithms are sensitive to magnitude. See: Milligan & Cooper (1988).

¹⁰⁸ Hierarchical clustering results, additional graphs, and statistics related to profile design are available at: <https://zenodo.org/record/8208858>

and cluster 2 (green) tended to be distinguished from cluster 3 (blue) by higher hedonism scores. Clusters 1 and 2 were largely split along the *Achievement/Power* values.¹⁰⁹ Because Schwartz's theory of values postulates that the order in which we prioritize our values should be conserved across contexts ("trans-situational") and clustering on general value scores provided more distinct clusters (de Wet et al., 2019; Schwartz, 2012), we selected these clusters for further profile development.

Three profiles were designed from the clusters by crafting personas (Rosson & Carroll, 2002), or creating narrative descriptions of each cluster based on their distinctive value characteristics (Figure 4-3). We used the Kruskal-Wallis test with a Dunn's post-hoc test to identify general values that were ranked significantly higher or lower between clusters ($p < 0.05$) (Dinno, 2015; Dunn, 1964; Kruskal & Wallis, 1952).¹¹⁰ There were no statistically significant differences of *Universalism* scores between the three profiles, which was ranked highly in every profile. We also considered the values that were ranked highest within each profile, excluding *Universalism*.

- 1) **Profile 1, Adventurer**, has significantly lower values scores for *Security*, *Tradition*, *Benevolence*, and *Conformity* compared to the other two clusters. Its highest ranked values are *Stimulation* and *Self-Direction*.
- 2) **Profile 2, Goal Setter**, has significantly higher *Power*, *Achievement*, and *Hedonism* scores than the other two clusters. These are also its highest ranked values.
- 3) **Profile 3, Helpful Neighbor**, has lower *Self-Direction*, *Stimulation*, and *Hedonism* compared to the other two clusters. Its highest-ranking values are *Benevolence*, *Security*, and *Conformity*.

¹⁰⁹ 2D graphs were also constructed but difficult to interpret because of the scale and overlapping points. They can be viewed at: <https://zenodo.org/record/8208858>

¹¹⁰ Some value datasets were normally distributed ($p > 0.05$), but others were not – hence the use of Kruskal-Wallis over One-way ANNOVA. Full statistics and additional graphs of each *value* are available at: <https://zenodo.org/record/8208858>




Figure 4-2: Plot of clusters according to *Power* (x-axis), *Achievement* (y-axis), and *Hedonism* (z-axis), where cluster 1 is red, 2 is green, and 3 is blue

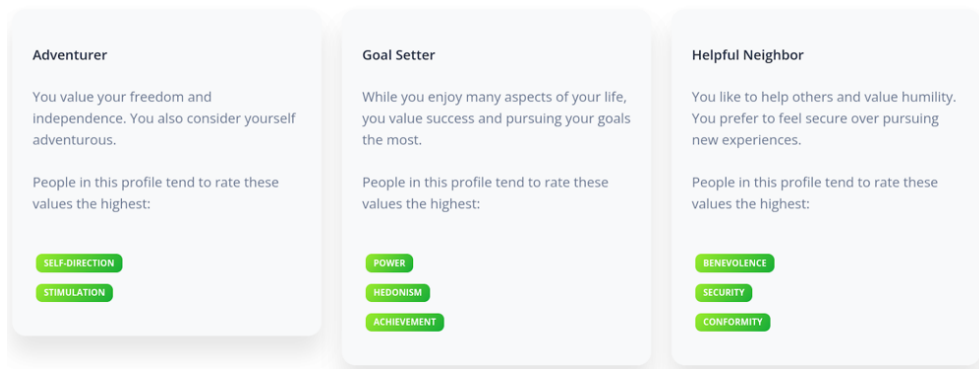


Figure 4-3: Presentation of VcPA profiles to participants on the Mock App Store

c VcPA Logics

The VcPA involves three primary features: selective notices based on a value profile; suggesting alternative applications that are more consistent with one’s selected profile; and exploratory notices to ensure the profile remains the best match (Table 3-2, Chapter 3). To accomplish this in the MAS, we firstly calculated a *minimal acceptability coefficient* using survey data in Phase I. The coefficient is the proportion of survey participants in the profile accepting the data collection practice required by the app that the least number of survey participants in that profile would be willing to accept. Selective notices (Figure 4-4) were triggered when an app’s practices did not match the profile, determined as a cutoff point of the above coefficient.¹¹¹ We began with a cutoff of 0.20 and tuned it by interacting with the

¹¹¹ Machine learning (Naïve Bayes, random forest, support-vector machine, k-nearest neighbors, and neural networks) was also tried to determine when a notice should be triggered based on the user’s selected profile and how well the app’s privacy permission requirements match privacy preferences associated with that profile. However, all accuracies were >60%.

MAS to ensure that notice frequency was not so high as to cause notice fatigue. A cut-off point of 0.1 was ultimately selected.¹¹²

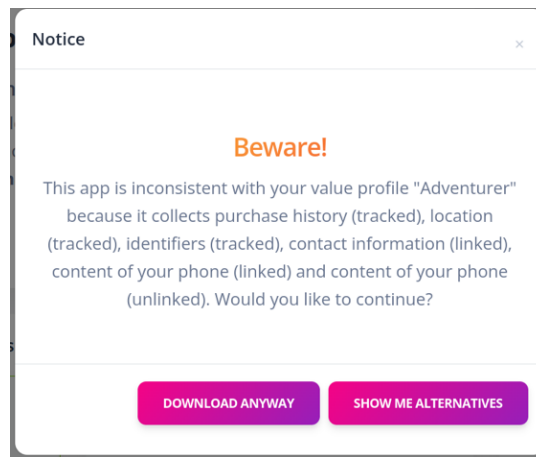


Figure 4-4: Selective notice example

On the MAS, we also added a “traffic light” system to the apps (Figure 4-1), where apps that were below 0.1 were “red” (and would trigger a notice if someone tried to download them); “yellow” if between 0.1 and 0.5; and “green” if above 0.5. Selective notices also included a button pointing to the “suggest alternatives” page, which included apps that matched the participants’ profile (coefficient>0.1) in the app’s family.¹¹³ If a participant ignored a selective notice, they were asked to specify a reason why.

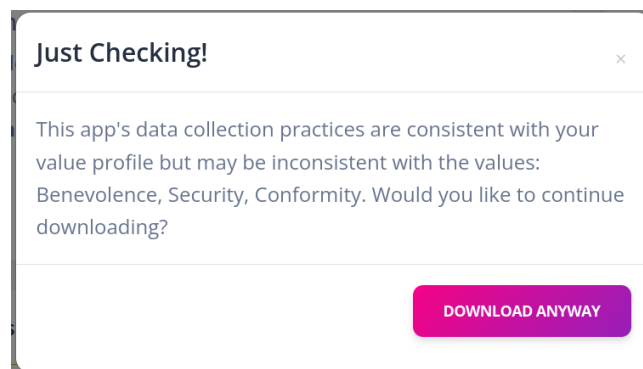


Figure 4-5: Exploratory notice example

¹¹² To demonstrate this further, consider a hypothetical health app Health Tracker, that asks for access to a user’s linked health and fitness information, unlinked location, and unlinked diagnostic data. A MAS participant indicates their desired profile as Goal Setter and then goes to “download” Health Tracker to their virtual smartphone. To calculate the coefficient, we use the number of survey participants in the selected profile who are willing to share the type of data asked for by the app. For our example, let us say 5% of survey participants who were in cluster 1 (Goal Setters) indicated that they were willing to share linked health information; 30% for unlinked location; and 50% for unlinked diagnostics. The minimal acceptability coefficient in the case of Health Tracker would be 5%, or 0.05. As the coefficient is less than the cutoff of 0.10, a selective notice would be triggered for the MAS participant when trying to “download” Health Tracker. All coefficients are listed in Appendix II.

¹¹³ See (a), previous.

The VcPA exploratory notices, utilized to check that the user's profile was still the best match (Figure 4-5), were also integrated into the Mock App Store. While it is unlikely that one's values would shift within a five-minute time, justifying a change in profile, we still wished to assess the exploratory notices in some manner that hopefully would not cause excessive notice fatigue. To accomplish this, notices were triggered between 3 minutes and 30 seconds and 4 minutes after the exercise began, when the participant clicked on an app.

d Entry and Exit Survey Designs

Before starting the exercise, participants were asked to take an entry survey with questions pertaining to basic demographic details, basic privacy attitudes, and values. Following the exercise, they were asked to take an exit survey about their views on the VcPA. The exit survey asked questions about the overall modality and perceived functionality of VcPA notifications and features (using a Likert Scale), as well as an open-ended question asking what could be improved. Both the entry and exit surveys contained questions asking participants to rank their level of overall privacy concern and smartphone privacy concern, and both were created on Microsoft Forms.¹¹⁴

e Feedback and Testing

In the same manner as the survey (see Section 4.2), feedback for the Mock App Store and VcPA were given in an iterative manner by collaborators, the PhD supervisory team, and the candidate's Graduate Research Committee (GRC).

4.3.3 Conducting the Study

After completing the entry survey, the participant was shown the profiles obtained from the clustering, along with a brief description and the top values associated with that profile (Figure 4-3). Participants were asked to select a profile that best matches their value set. They were then prompted to browse the store and "download" apps they would be interested in having on their virtual "smartphone."¹¹⁵ App icons and detailed descriptions were presented to enable participants to select relevant apps to download to their "virtual smartphone," with VcPA elements (notices and the "suggest alternatives" page) appearing to help users where applicable. They were also given the option to remove apps from their "virtual smartphone" at any time during the exercise. Participants had five minutes to complete this task before being re-directed to the exit survey.

a Participant Recruitment and Selection

The Mock App Store Study was conducted with a different smaller group of participants, recruited separately from the general survey (Phase I). Based on existing personal privacy assistant literature (Liu et al., 2016), we aimed to recruit roughly 100 participants. Participants were recruited via snowball sampling over researcher's professional networks in a similar manner to the survey (Section 4.2). Like the survey, all participants were required to be current or previous smartphone users who were 18 years or older and possess a native or fluent command of English.

¹¹⁴ Entry and exit survey available in Appendices III and IV, respectively.

¹¹⁵ Demo video of the MAS with VcPA available at: <https://youtu.be/ziGoowteN6E>

Consent for the entire experiment (Entry Survey, Mock App Store exercise, and Exit Survey) was asked at the time of recruitment. An information sheet was also included at the time of recruitment with the option to withdraw or contact the researcher for additional information. Participants were lastly asked if they would be willing to be re-contacted for an optional follow-up interview (detailed in the upcoming section, Section 4.4).

b Data Collection and Analysis

Each interaction with the MAS, such as which profiles were selected, which apps were downloaded or removed, and interactions with selective and exploratory notices, was recorded at the end of the exercise. Where applicable, interaction differences between the three profiles were analyzed using t-tests. Due to unequal sample sizes, data from the entry and exit survey were analyzed using a two-sample unequal t-test. Logs from the MAS that did not include any downloaded apps were excluded.

Section 4.4 Phase 3: Post-Study Interviews

4.4.1 Phase III Summary

The final qualitative phase consists of semi-structured interviews. Given the richness and complexity of human values, the aim of these interviews was to gain a deeper understanding of how user values affect privacy and app choices to complement data collected in Phase I and II (Table 4-1). Questions were designed to elicit affective responses to ascertain participants' values as well as elicit more comprehensive feedback on the VcPA. Analysis was done using a process of Reflexive Thematic Analysis to identify values, tensions, and feedback in a predominately latent and inductive manner using a critical realist paradigm (Braun & Clarke, 2022). Interviews were conducted between one and four months after the participant completed the Mock App Store Study (July-October 2022).

4.4.2 General Interview Structure and Design

Before the interview, participants were asked to answer a quick survey asking them for their demographic details. Demographic information from the short survey was reported in aggregate for descriptive statistics purposes only (e.g., "7 of participants were male and 8 female"). Participants also received an information sheet and were asked to complete a second consent form. The interviews themselves were 22-59 minutes in length and conducted in English. They were conducted over Microsoft Teams and recorded. To ensure participant privacy, participants were informed they could join using Google Chrome in Incognito mode. The interviews were semi-structured to allow participants sufficient space to elaborate on their privacy and app decision-making process while still staying relevant to the research questions. Questions were drawn from a question bank¹¹⁶ but were flexible depending on the content of the interview. To ease into the interview, questions about the VcPA were discussed first before opening to a discussion of the participant's real life data privacy or smartphone interactions.

¹¹⁶ The question bank is in Appendix IIX.

4.4.3 Question Design

Questions aimed to capture profile understanding, feature reception, and app selection process in the MAS as well as eliciting personal values and value tension when choosing an app or making a data privacy decision. The original interview questions were piloted with a group of colleagues at University of Galway (UG) and in other Irish higher education institutions, as well as a few family members. These were further modified to reflect feedback and advice following completion of a graduate methodology course (2021-2022), such as modifying or removing leading questions. In addition, questions concerning values were modified to probe participant's commitments as a representation of values instead of asking directly about values. This is based on the 4DT understanding of values described in Chapter 3 (see Figure 3-1), where values can be understood as clusters of commitments we take upon ourselves as self-governing agents.¹¹⁷ Critically, based on this understanding, people may hold the same value *a*, but the cluster of commitments pertaining to *a* may be different. This could lead to different understandings of what *a* is. Compounding this, linguistic variation may result in different understandings, definitions, or words for value *a* between individuals. To account for this diversity and decrease participant confusion, questions were designed to extract motivations/commitments, areas of value tension (that may be *double binds*) between values (Tables 3-1 and 3-2, Chapter 3), and any relevant affective responses from which the latent value could be inferred. This has implications for the analysis, which must consider this large role of the researcher in knowledge production (discussed in greater detail in Section 4.4.5).

4.4.4 Participant Recruitment

During recruitment for the Mock App Store (MAS) exercise, participants were asked if they would be willing to be recontacted for a follow-up interview. Based on existing Personalized Privacy Assistant (PPA) literature (Colnago et al., 2020), we aimed to have around 15-20 participants, considering also both saturation (no new content) and informational power (sufficient richness of data) (Braun & Clarke, 2022; Glaser & Strauss, 1999; Malterud et al., 2016; Sandelowski, 1995). Before the interview, participants were sent an email with photos and an explanation of the MAS exercise to jog their memory. They were also shown the MAS again during the interview itself. Questions about the exercise were used as a starting point for discussion before moving on to additional questions from the question bank concerning **RQ1**.¹¹⁸

4.4.5 Data Analysis and Coding

Because the interview qualitative data was meant to complement Phase I and Phase II quantitative data through a process of convergent design (Fetters et al., 2013), thematic analysis (TA) was selected for ease of comparison of qualitative data with quantitative data. For our purposes, TA allowed us to explore relevant values by grouping codes of

¹¹⁷ In brief, we can understand values as clusters of commitments we take on as self-governing agents with a positive orientation towards some shared state or object (Killmister, 2017). These commitments are understood to be generated from certain beliefs, goals, and attitudes about who we are and how we want to exist in the world. In this case, data privacy choices are a decision point that can be understood as a fulfillment of or violation of one's values, with the decision whether to download an app representing one such privacy decision point (in addition to, for example, more traditionally explored choices as engaging with cookie consent notices online).

¹¹⁸ See footnote 116.

commitment statements or affective responses into a theme for each identified value. Themes were also developed based on coded VcPA feedback in the interview. These themes could then be compared with the value-privacy relationships identified in the survey data.

The process of TA is described more in the following subsections. In (a), the theoretical considerations of a particular flavor of TA that we utilized, Reflexive TA (Braun & Clarke, 2022), is explored. This is followed by a breakdown of the steps of Reflexive TA as they were applied to this study.

a Theoretical Considerations of Reflexive Thematic Analysis

Specifically *Reflexive* Thematic Analysis (Reflexive TA), as described in Braun and Clarke (2022), was selected to capture my (the researcher's) role in knowledge production. Reflexivity refers to a thoughtful, reflective researcher constantly owning their subjectivity and role in knowledge production. In this case, reflexivity was crucial due to the personal, abstract nature of values, and a recognition of my own motivations for undertaking this work that will inevitably affect my interpretation of the data. In addition, I played an active role in interpreting what values were latently represented by particular participant statements and selecting a method like Reflexive TA allowed critical space for me to reflect upon my own positionality when interpreting their statements.

(1) Ontological and Epistemological Considerations

Different qualitative approaches are amenable to Reflexive TA. Experimental TA involves the researcher to take *hermeneutics of empathy* orientation – interpreting data by staying close to the participant's words and understandings of the world (Braun & Clarke, 2022). Critical TA, in contrast, utilizes an orientation of *hermeneutics of suspicion*, that is, seeking to interrogate meaning *behind* the words. In this case, a primarily critical qualitative approach was taken to answer the two research questions (Table 4-1). For example, if a participant expresses that they find targeted advertising creepy, the values behind that (why *might* they feel this way?) were inferred instead of directly taken from the participant's own words. This interpretative process is described in greater detail later in this section when discussing the coding process.

In addition, to capture the interplay between coding for what is said in the data and the meaning behind it, the investigation was situated within the larger paradigm of *critical realism* (Braun & Clarke, 2022). This takes a contextualist epistemological standpoint that recognizes that many representations of reality are possible because people's interpretations of reality are influenced by their own unique positionality and contexts, and that researchers cannot remove people from these contexts and study them in a vacuum. Qualitative data is here viewed as subjective but still able to reveal underlying meaning and phenomena. Participants in the study are acknowledged to be able to construct their own meaning and understanding of values when making privacy decisions, but the researcher is still able to extract some underlying meaning concerning values, app choices, and privacy preferences.

Writing this critical realist paradigm, coding of values utilized a mixed inductive and deductive approach. Deductively, coding was driven by the specific understanding of values, theoretically grounded in the understanding of values in Chapter 3. This specific understanding of values helped inform what constitutes a "value," understanding values as a cluster of commitments that could also be in tension with each other and, in cases of equal valuing, a *double bind* situation. An inductive coding approach – without the

theoretical mapping of 4DT – was also utilized, focusing on the participant’s particular understanding of areas of value tension between values. While still theoretically aware of 4DT its theoretical implications for promoting or undermining value-consistent privacy decisions, it was also important to loosen our grip on this theory to identify other areas where participants felt at odds with their values, especially when it came to value tensions. In addition, while 4DT was the model of autonomy that was utilized to further translate the idea of value-centered privacy into a tangible privacy assistant (Chapter 3), this does not mean that it is necessarily the catch-all for conceptualizing all the ways value-centered privacy choices could manifest themselves. Opening to a level of inductive coding allowed us to identify areas of value tension or understandings and insights into values that may not be captured by 4DT. It also provided us with insights into the strengths and limitations of the 4DT or value-centered approach.

In addition, coding utilized both semantic and latent approaches to meaning. A more semantic (word for word; “direct” meaning) approach was initially used when coding participant statements, with a latent (“patterns” of meaning) utilized when forming values (themes) and during the second pass through the data. This was because participants rarely talk about their values as such (see Section 4.4.3). Affective responses helped inform which latent value could be inferred when coding. Statements pertaining to the VcPA were coded using a semantic approach, staying close to the explicit words of participants and their experience with the VcPA to identify areas of improvement and provide further insight into the strength and limitations of a value-centered approach.

b Conducting Reflexive Thematic Analysis

(1) Overview of Reflexive Thematic Analysis Process

Thematic analysis involves coding a text to identify themes. To further interrogate the meaning, researchers’ conducting Reflexive TA also keep a reflexive journal throughout the entire study and analysis. After data collection and transcription, Reflexive TA involves six phases, which are iterative and non-linear: 1.) familiarizing oneself with the data; 2.) coding; 3.) generating initial themes; 4.) developing and reviewing themes; 5.) refining, defining, and naming themes; and 6.) writing up results. In this case, familiarization occurred concurrently with conducting the interviews and transcribing them, and the coding and theme phases occurred iteratively. Familiarization and transcription were conducted from fall 2022-winter 2023, with coding and theme development finalized by the summer 2023. Each phase as it pertains to this investigation is described in greater detail in (2-4).

(2) Keeping a Reflexive Journal

Besides the distinct phases of Reflexive TA, it is also critical that an ongoing journal is kept that reflects upon the researcher’s own personal situation (values, beliefs, motivations, and attitudes) to the work, as well as functional (research design) and disciplinary influences. In this work, I kept a journal to reflect upon my own positionality as the researcher conducting the interviews and the analysis. After an initial brainstorm concerning my personal, functional, and disciplinary influences, the journal was continually updated throughout the interviewing and analysis processes to raise awareness of my positionality and role in meaning-making, especially times when the interview process or participant statements caused feelings of concern, anxiety, or disagreement.

(3) *Transcription and Familiarization*

After each interview, I noted any initial points of interest in my research journal. Otter.AI¹¹⁹ was used to generate the initial transcripts which I then checked against the recording for accuracy. Any possibly identifying information (such as reference to a place of work, city, etc.) were also removed. Interviews were transcribed clean verbatim – fillers words such as “um” and laughter were not transcribed unless I deemed it necessary for understanding the text (e.g., capturing affective responses). I also continued familiarizing myself with the data by noting any additional thoughts for each interview and thoughts pertaining to the dataset as a whole while transcribing. I reviewed these notes in-depth before beginning the coding process.

(4) *Iterative Coding and Theme Development*

The interviews were coded in NVivo. Coding and theme development occurred in an iterative, concurrent manner. Codes are an analytical tool that captures a group of related researcher insights from engaging with the data. Themes capture shared meanings – or central organizing concepts – between codes. Initial theme development also occurred while coding to explore possible clustering of codes. In this case, themes were values, encompassed VcPA feedback, or captured another related interesting concept in the data (such as feelings regarding tech companies). I completed this iterative coding and thematic process twice.

Codes and themes were also cleaned and refined in an iterative manner while coding, with dedicated time spent to code and theme refinement both between passes through the data and after the final pass. Before coding each participant, I reviewed my notes for that participant I had taken during the familiarization phase (3). During coding, codes were reviewed, refined, and merged throughout the coding process to be appropriately broad and representative. To refine codes, codes were checked using the “take away the data” test (Braun and Clarke, pg. 71) to verify that the code accurately captures the content and my analytical take. Codes were also checked for consistency and against other similar codes to ensure codes were sufficiently distinct. Themes were also checked to ensure that the data contained in each was not too diverse and had a coherent central organizing concept. Themes were also checked against each other to ensure the theme boundaries were clear. Themes were merged or split when deemed appropriate. In particular, initial themes were grouped together following the first pass to better capture some shared commonalities between them. During the second pass, special attention was paid to ensure that the codes and themes matched the source material. Interview transcripts were also shuffled for the second pass to ensure an evenly coded dataset.

(5) *Writing Up Results*

The final stage of Reflexive TA is writing up results, which involves weaving the themes into a coherent story. Results in this case were interwoven with the results of Phase I (Chapter 5) to answer **RQ1** and Phase II (Chapter 6) to answer **RQ2** (Table 4-1). In the next chapter, I will present relevant interview results alongside Phase I survey results to answer **RQ1: What is the relationship between values and privacy preferences when deciding to download an app, if any?**

¹¹⁹ <https://otter.ai>

Chapter 5 How Do We Value Data Privacy?

*Stunning 8K-resolution meditation app
In honor of the revolution, it's half off at the GAP
[...]
Female Colonel Sanders, easy answers, civil war
The whole world at your fingertips, the ocean at your door
[...]
Full agoraphobic, losing focus, cover blown
A book on getting better hand-delivered by a drone
Total disassociation, fully out your mind
Googling "derealization," hating what you find
[...]
There it is again, that funny feeling
That funny feeling*

Bo Burnham ("That Funny Feeling")

Section 5.1 Chapter Overview

This chapter presents results from the online value and privacy preference survey (Phase I) and the relevant interview results (Phase III) to answer **RQ1**: *What is the relationship between values and privacy preferences when deciding to download an app, if any?* In particular, we explore *how we value privacy*: how values, privacy preferences, and app choices correlate, are understood, and are conceptualized by users. Answering **RQ1** provides us with empirical insights into how values are involved in privacy decision-making as a means of promoting more value-centered choice. It also allows us to evaluate the 4DT understanding of the role of values in privacy decisions, where values are understood as clusters of commitments that inform us to act.

To these ends, the survey provides quantitative insights into the correlations between the 10 general, life-guiding Schwartz values (hereafter just *Value*, e.g., *Hedonism*); the relevance of these 10 values in the context of deciding whether to download an app (*Value App*, e.g., *Hedonism App*); the difference between values as general, life-guiding principles and the values in the context of downloading an app (*Value Difference*, e.g., *Hedonism Difference*); the app under consideration (Lose It! or OpenLitterMap, *App*); and 40 different Apple privacy preferences (e.g., *unlinked location*). In summary, we observed many weak ($\rho < 0.5$) correlations between privacy preferences and values, with correlations and relevant values varying depending on which app was being considered. This suggests that values are involved in a primarily context-dependent (here, app-dependent) manner. Observed methodological limitations of using Schwartz values and Apple privacy preferences in our survey suggest that these results will need to be further complemented with research using alternative methodologies.

To this end, the quantitative results from the survey data are further contextualized with results from semi-structured interviews. These interviews provide a deeper, richer understanding of how we value data privacy – including how we describe, present, and understand the role of our values and their tensions. From the interviews, we identified six major value themes that were motivationally related to privacy decisions in a manner driven by both an individual's understanding of the value and the context in question.

How Do We Value Data Privacy?

Similarly, tensions between values were also linked to context and individual interpretations of the values. Despite value relevance to privacy or app decisions mostly varying based on the individual and the context, a handful of values such as *Use*, *Community*, and especially *Control* were frequent and often in tension with each other. These results were largely consistent with the 4DT-informed understanding of value-centered privacy decisions and values as clusters of commitments. We identified expected phenomena such as instances of *weakness of will*, the “apathetic user,” and inability to *self-realize* due to lack of relevant privacy controls or alternative choices. We also found that interview participants reported many strategies to resolve tensions between values, although these resolution strategies could be better supported. This suggests that areas of value tension where current resolution strategies are unsatisfactory could also be improved by a value-centered approach. However, the concept of *double binds* allowed us to identify the areas in *most* need of a value-centered intervention. Although the line between value tensions that can be resolved and *double binds* existed on a spectrum within the data, the presence of structural factors (e.g., social media monopolies or surveillance capitalism) were defining features of the tensions that *most* resembled *double binds* (most “*double bind*-like). Other interesting observations from the interviews, including the role of privacy regulation and participant views on data-based business models, are also briefly touched upon.

5.1.1 Collaborator Contributions

The studies presented in this chapter were conducted in collaboration with Prof. Dr. Mathieu d’Aquin (supervision guidance, data analysis, manuscript feedback), Dr. Heike Felzmann (supervision guidance and manuscript feedback), Prof. Dr. Kathryn Cormican (supervision guidance and manuscript feedback), Dr. Dave Lewis (supervision guidance), Dr. Ilaria Tiddi (manuscript feedback), and Dr. Dayana Spagnuolo (data analysis). I (the PhD candidate) designed and conducted the study, as well as worked with collaborators at all stages of data collection, data analysis, and results write-up.

5.1.2 Relevant Papers and Conference Contributions

Some material in this chapter, including certain text and figures, has been previously published or presented in the following:

Carter, Sarah E., & Felzmann, Heike. (2023). How do we value data privacy? Insights and design implications. To be published in: *Engineering and Value Change* (part of: *Springer Philosophy of Engineering and Technology series*). Abstract available at: <https://zenodo.org/record/8367542>

Carter, S.E., d’Aquin, M., Spagnuolo, D., Tiddi, I., Felzmann, H., & Cormican K. (2023). The privacy-value-app relationship and the value-centered privacy assistant. ArXiv. <https://arxiv.org/abs/2308.05700>

Carter, Sarah E., & Felzmann, Heike. (2023, April 21). How do we value data privacy? Initial results from semi-structured interviews. Forum on Philosophy, Engineering, and Technology (fPET2023), Delft, the Netherlands. Zenodo. <https://doi.org/10.5281/zenodo.8204406>

Section 5.2 Online Value and Privacy Preference Survey

5.2.1 Section Overview

To start exploring the relationship between our values, our privacy preferences, and our app choices (**RQ1**), we first conducted an online survey of 273¹²⁰ smartphone users' values and privacy preferences when considering whether to download one of two apps (Lose It! and OpenLitterMap). Our results suggest that values and privacy preferences are related in a primarily app or context-dependent manner. We found 215 weak ($\rho < 0.5$) but potentially interesting correlations between privacy preferences, values, and which app was being considered. The strongest correlations occurred between which app was being considered and a user's values (for example: OpenLitterMap and *Universalism*). When looking at the value-privacy preference relationship within each app, 197 additional correlations were observed, and different values were considered more relevant depending on the app being considered. Taken together, the multitude of value-privacy preferences relationships suggests that values are related to privacy preferences, and therefore privacy decision-making.¹²¹ However, some of these results also suggest limitations of the method used, supporting the need for further contextualization with the interview results.

a Participant Demographics

In total, we obtained 305 engagements with the online value and privacy preference survey. The survey automatically stopped for those who did not meet our criteria or did not complete the consent form. This resulted in 273 complete and usable responses. Participants were randomly assigned to either the Lose It! and OpenLitterMap version of the survey and were split fairly evenly between the two (147 vs. 126) (Appendix VI).¹²² The majority (168) of participants identified as women, 95 as men, and 10 as other/non-binary/prefer not to say. Age was heavily dominated by adults (ages 25-64), with 204 participants, and 62 participants were young adults (18-24); only 7 were older adults (65+). The majority also had (or were in the process of obtaining) a doctoral or master's degree (214), with 50 for bachelor's and 8 for a secondary degree. Nationalities were grouped by continent, with the majority having European nationalities (176), followed by North America (41), Asia (38), and 18 other/prefer not to say. There were 157 fluent and 116 native English speakers, and nearly all participants (271) currently had a smartphone.

¹²⁰ There was a total of 305 engagements with the survey, but participants were automatically exited from the survey if they did not consent or meet the relevant criteria (smartphone use and English proficiency). See: <https://zenodo.org/record/8208858>

¹²¹ Recall that our privacy preference is understood from a 4DT lens as the result of our deliberation on our values – that is, what we best *ought* to do. See Section 3.2.

¹²² While members of each demographic group were spread fairly evenly between the Lose It! and OpenLitterMap versions of the survey, there were a few notable differences. In terms of nationality, there was a greater percentage of Asian nationalities for LoseIt! (17%) than OpenLitterMap (10%) and greater percentage of European nationalities for OpenLitterMap (68%) than LoseIt! (61%). In terms of age group, there was a greater percentage of adults who took the LoseIt! (80%) version than the OpenLitterMap (69%), and a greater percentage of young adults of OpenLitterMap participants (28%) than for LoseIt! (18%).

5.2.2 Majority of Participants Believed Values are Involved in App Choice

When asked, most survey participants (73%) believed that their values were a factor when choosing a smartphone app, although perhaps not all the time: 30% of participants answered “yes” and 43% answered “sometimes.” Only 15% said “no” and 11% were unsure. This could be because participants were aware of the purpose of the survey. However, this still suggests that most participants (minus the 15% who said a definite “no”) are at least open to the idea that values could or should contribute to app choice. While it is often hard to consciously pinpoint the motivation and effects of values on our behavior due to their abstract nature, it is promising that users were open to the idea that values could or should be involved in their decision-making. This suggests that they may also be open to a value-centered privacy approach or intervention.

5.2.3 Values, Privacy Preferences, and Apps are Weakly Correlated and Correlations are Context-Dependent

From the survey data, some correlation also appears to exist between values and privacy preferences (Figure 5-2). There were multiple significant, weak ($\rho < 0.5$) correlations between general values and overall app privacy preferences. When survey results for both versions were grouped together, there were 215 significant ($p < 0.05$) and weak ($\rho < 0.5$) correlations between *Value*, *Value App*, *Value Difference*, *App*, and privacy preferences, where *Value* is the general value score (as a life-guiding principle, derived from the TBHV); *Value App* is the value score when considering whether to download either LoseIt! or OpenLitterMap; *Value Difference* is the difference between *Value* and *Value App*; *App* is either LoseIt! or OpenLitterMap; and the privacy preferences were based on Apple Privacy labels (e.g., *tracked location* data) (Table 5-1).¹²³ The strongest overall correlations were between *App* and *Universalism App* and *Benevolence App*, suggesting that the app being considered has the strongest correlation to a value’s relevance (> 0.35). *Universalism Difference* and *Benevolence Difference* were also inversely correlated with each other between Lose It! and OpenLitterMap (both 0.32 for Lose It!, and -0.32 for OpenLitterMap). These *Value App* correlations make sense given the functions of Lose It! and OpenLitterMap and suggests a context-dependent relationship between values and privacy preferences. This is discussed in (a), where *Value App* scores between the two apps are compared in greater detail.

Besides correlations with the apps, there were a few global privacy preference-app relationships of note (Figure 5-2). We also suggest tentative interpretations of these correlations. The value/privacy preference correlations were: 1) *no unlinked data* was negatively correlated to *Hedonism App* and *Achievement App* (both -0.23); and 2) *unlinked location* was positively correlated with *Universalism App* and *Benevolence App* and negatively correlated with *Universalism Difference* (0.25, 0.24, -0.24). The first suggests that those who value *Hedonism* or *Achievement* may be more likely to be willing to share unlinked data with the app. This makes some sense when considering the goal of Lose It! to meet one’s health and fitness goals (*Achievement*); a user may be willing to share more data with the app to meet their goals. The reason behind the *Hedonism App* correlation is less clear but could be because increased data sharing is associated with increased app engagement and more fun. The second grouping of correlations concerning *unlinked location* data suggests that higher levels of *Universalism* and *Benevolence*, as well as the

¹²³ Recall from Chapter 4 that the survey was designed using the Short Schwartz Value Survey (SSVS) and Apple Privacy Labels. The survey is included in Appendix I.

How Do We Value Data Privacy?

degree that one's *Universalism* levels when engaging with an app differ from their baseline, increase one's willingness to share their location data with the app. The reason behind this is not intuitive but could relate to a willingness to share data to better the world, as is the case with the other app (OpenLitterMap), where the location would need to be shared to identify where the litter is.

While interpretations of the individual privacy preference-*Value/Value App/Value-Difference* are difficult to make, these results, taken together, suggests that *some* value-privacy preference relationships exist independent of the app context.

a Lose It! and OpenLitterMap Specific Correlations

Because which app was being considered (Lose It! or OpenLitterMap) correlated the strongest with *Value App* and *Value Difference* scores, correlations were also calculated between *Value*, *Value App*, and privacy preference for the Lose It! and OpenLitterMap datasets separately (Appendix VII). From these analyses, we identified 197 additional weak and significant value (app)/privacy preference correlations that are app-specific.¹²⁴ These results suggest that values and privacy preferences are primarily related in the context of an app. We also aimed to initially intuitively understand why some values and privacy preferences may be related to either Lose It! or OpenLitterMap based on the function of the app.

(1) *Lose It!*

For Lose It!, all correlations were weak ($\rho < 0.5$), with 120 significant correlations identified. However, some notable patterns were observed, with preliminary interpretations presented here.

Firstly, *unlinked phone content* was weakly correlated with almost every *Value/Value App/Value Difference* score except *Hedonism Difference*, *Hedonism*, *Tradition*, *Benevolence*, and *Achievement Difference*. Similarly, *tracking contact information* and *tracking finance information* were correlated with most *Value/Value App/Value Difference* scores. This suggests that these privacy preferences could be closely intertwined with values when it comes to Lose It!, whether it be general, life-guiding values (*Values*) or values when choosing an app (*Value App*).

Secondly, *Hedonism App* and *Hedonism Difference* also had many weak correlations with privacy preferences. The strongest pairs were *linked general diagnostic data* (-0.28 and -0.29), *linked sensitive information* (-0.20, -0.25), *tracked general diagnostic data* (0.18 and -0.26), and *unlinked browsing history* (0.23, -0.22). This suggests the importance of the value, *Hedonism*, when deciding whether to download Lose It! (*Hedonism App*), and how closely that matches one's overall ranking for *Hedonism* (*Hedonism Difference*), may have small influences on these privacy preferences. In addition, this supports what was suggested in the previous section that one may be willing to share more data to increase engagement and pleasure from app use. This sharing, however, has limits, based on the negative correlations between *Hedonism App* and *linked diagnostic data* and *linked sensitive information*.

Thirdly, *Tradition* is negatively correlated with desiring *no tracking data* (-0.25) and *no linked data* (-0.18), suggesting that those who value *Tradition* less may be slightly

¹²⁴ It is also important to note that the participant demographic between the Lose It! and OpenLitterMap versions of the survey were similar (Appendix VI). Considering that the ratios in the demographics are close, this suggests that, in the case of this study, that any differences we observe between the values in the Lose It!/OpenLitterMap context is due to the app itself and not demographic differences.

How Do We Value Data Privacy?

more likely to be comfortable sharing linked and tracking data with Lose It! *Tradition*, again defined as “being humble, modest, and respecting tradition” (Figure 2-2), could be tied up in abiding by more traditional norms and values around privacy and sharing one’s life with the world.

Lastly, *Stimulation App* and *Stimulation Difference* were correlated with *linked contact information* (0.24 and -0.21). This suggests that how relevant a participant finds *Stimulation* when deciding whether to download Lose It!, and how much this differs from their overall value of *Stimulation*, may be tied to their willingness to share contact information. One possible explanation could be the increased stimulative value of contact from apps through emails, pop-ups, or other notifications.

(2) *OpenLitterMap*

Within the OpenLitterMap dataset, correlations between *Value*, *Value App*, and privacy preferences were again weak, and patterns harder to identify than with Lose It!. There were also less significant correlations than Lose It! (77), although the strongest correlation was slightly stronger than Lose It! (-0.30 vs. -0.29). Like Lose It!, *Stimulation App*, *Stimulation Difference*, and *Tradition* showed weak correlations with (different) privacy preferences: *no unlinked data* for *Stimulation App* (-0.25) and *Stimulation Difference* (0.25) and *unlinked health and fitness information* for *Tradition* (0.25). Also notable was *Power App* and *unlinked health and fitness information* (0.28) and *linked location* (0.26), suggesting that those who value *Power* more when engaging with OpenLitterMap were more likely to give away these types of data. There were also some interesting correlations between *no linked data* and *no tracking data*, including *Conformity* and *no linked data* (0.23) and *Self-Direction Difference* and *no tracking data* (0.23). This suggests that the more one values *Conformity* as a life-guiding principle, the more likely they are to not share linked data when engaging with OpenLitterMap.

Most interesting, however, were the relatively high correlations between *unlinked location* and *Universalism Difference* (-0.30) and *Universalism App* (0.28). The first suggests that the more one’s value of *Universalism* changes when considering downloading OpenLitterMap (*Universalism Difference*) from their baseline *Universalism* levels, the more their willingness to share *unlinked location* data also increases. The second correlation suggests that those who rank *Universalism* higher when considering the app (*Universalism App*) are also more likely to find sharing *unlinked location* data acceptable. This makes sense given the goals of OpenLitterMap to help the environment, and this context dependence of values is discussed more in the next section.

In summary, while the individual correlations are again quite difficult to interpret, the existence of many distinct value-privacy preference correlations for Lose It! and OpenLitterMap suggests that values and privacy preferences are primarily related in a context or app-specific manner.

How Do We Value Data Privacy?

Table 5-1: Features analyzed in survey data¹²⁵

Value¹²⁶	Value App	Value Difference	App¹²⁷	Privacy Preferences¹²⁸
<i>Definition: universal, life-guiding principle</i>	<i>Definition: value relevance when considering whether to download an app</i>	<i>Definition: Difference between Value and Value App</i>		<p><i>Tracking: A form of linked data that is further combined with data from other sources using identifiers</i></p> <p><i>Linked: data that is connected to a user's name or other identifiers.</i></p> <p><i>Unlinked: data that has been stripped of a user's name and other identifiers (e.g., phone's unique ID)</i></p> <p><i>Format: linkage + data type (e.g., unlinked browsing history)</i></p>
				Data Types
<i>Benevolence</i>	<i>Benevolence App</i>	<i>Benevolence Difference</i>	Lose It!	<i>Health and fitness information</i>
<i>Universalism</i>	<i>Universalism App</i>	<i>Universalism Difference</i>		<i>Financial information</i>
<i>Hedonism</i>	<i>Hedonism App</i>	<i>Hedonism Difference</i>		<i>Location</i>
<i>Achievement</i>	<i>Achievement App</i>	<i>Achievement Difference</i>	OpenLitter Map	<i>Sensitive information</i>
<i>Stimulation</i>	<i>Stimulation App</i>	<i>Stimulation Difference</i>		<i>Contacts</i>
<i>Power</i>	<i>Power App</i>	<i>Power Difference</i>		<i>Phone content</i>
<i>Conformity</i>	<i>Conformity App</i>	<i>Conformity Difference</i>		<i>Browsing history</i>
<i>Tradition</i>	<i>Tradition App</i>	<i>Tradition Difference</i>		<i>Search history</i>
<i>Security</i>	<i>Security App</i>	<i>Security Difference</i>		<i>Purchase history</i>
<i>Self-Direction</i>	<i>Self-Direction App</i>	<i>Self-Direction Difference</i>		<i>Usage data</i>
				<i>Diagnostic data</i>
				<i>Contact information (linked and tracking only)</i>
				<i>Other identifiers (linked and tracking only)</i>
				<i>None (no data collected)</i>

¹²⁵ See footnote 123.

¹²⁶ Definitions for each value are presented in Figure 2-2. The theory behind these values (the TBHV) is presented in Section 2.4.

¹²⁷ App links are in footnote 97.

¹²⁸ The Apple Privacy Preferences options are presented in greater detail in Table 4-2.

How Do We Value Data Privacy?




Figure 5-1: Heatmap of significant Spearman correlations ($p < 0.5$) for: total dataset *Value*, *Value App*, *App*, and privacy preference

5.2.4 Different Values were More Relevant Based on App

Because app-specific value-privacy preference correlations were observed, we next identified which values were more applicable to Lose It! and OpenLitterMap. To do this, we looked at the differences between *Value App* scores for Lose It! and OpenLitterMap. The importance of values was indeed different depending on which app the participant was considering. *Value App* scores were significantly different for all values except *Tradition App*, *Conformity App*, and *Security App* (Figure 5-2).¹²⁹ This suggests that the observed differences in the value-privacy preference relationship may be due to different values being relevant depending on the context in question.

¹²⁹ It is also notable that *Value* scores between the LoseIt! and OpenLitterMap cohorts only differed significantly for *Tradition*. When normalizing *Tradition App* by *Tradition*, we found significant differences between LoseIt! and OpenLitterMap *Tradition App* (towards Lose It!). This suggests that the lack of significance for *Tradition App* between apps may be due to how much a participant valued *Tradition* as a general life-guiding principle.

How Do We Value Data Privacy?




Figure 5-2: Differences in *Value App* ($p < 0.05$, unequal t-test) by t-statistic, where a positive value indicates higher importance to Lose It! and negative to OpenLitterMap

Many of these differences between *Value App* scores make intuitive sense based on the app function. Indeed, previous studies have suggested that the values vary by type or app or the app function (Liu et al., 2016; Nurwidyantoro et al., 2022) – something that is supported by our results. When investigating GitHub issues for three open-source Android applications, Signal, K-9, and Focus, Nurwidyantoro and colleagues (2022) observed a difference of relevant values when mapping observed Schwartz values between the apps.¹³⁰ They postulated that the difference they observed made sense based on the function of the app – for example, the identified *sense of belonging* (grouped under the Schwartz value of *Security*) with Signal, which involves connecting with others and messaging. Similarly, in our results, the values appear to match with the app’s function. *Universalism* (“the welfare of our broader world,” Figure 2-2) and *Benevolence* (“helping others”) are relevant when considering an environmental app like OpenLitterMap (focused on creating open-sourced litter datasets to work towards a cleaner world). *Achievement* (“success”), *Self-Direction* (pursuing goals), and *Stimulation* (“living a varied and challenging life”) also intuitively align well with an exercise and diet app like Lose It!.

Other differences between Lose It! and OpenLitterMap *Value App* scores, however, were more surprising, and do not intuitively match with their functionality. For example, *Power App* and *Hedonism App* were rated higher for Lose It!. *Hedonism* (“self-indulgence, pleasure”) seems contrary to exercise and weight loss goals, and indeed, *Hedonism* levels have previously been reported as negatively related to health app use (Mejova & Kalimeri, 2019). However, unlike Mejova and Kalimeri (2019), we were not asking what values cause you to download Lose It!, but the importance of each value in your decision whether to download Lose It!.¹³¹ It could be that someone who values *Hedonism* highly does not want Lose It! and finds it important in their decision *not* to download it. Regardless, just

¹³⁰ Interestingly, they also found that app value statements – for example, stressing the value of privacy for Signal and Focus – also seemed to influence the values popping up on the GitHub discussions. I return to this point when discussing interview results in Section 5.3, which suggests that these value statements are also influential in choosing value-consistent apps.

¹³¹ When asked whether they would download the app (Lose It or OpenLitterMap) after a brief description, most participants would not download the app. For Lose It!, 62 people said they would not download it; 32 would; and 53 were unsure. For OpenLitterMap, 62 people said they would not download it; 21 said they would; and 43 said they were unsure.

How Do We Value Data Privacy?

because the relationship does not intuitively match the function, does not suggest that a relationship between value and function does not exist. More investigation is needed to tease out what exactly the relationship may be by utilizing qualitative investigations to discuss one's motivation for downloading or not downloading apps in more depth. Our analysis of the interviews provides initial insights into these relationships (to be discussed in Section 5.3), allowing us to look more deeply at the motivations and thought processes behind the data privacy choices users make.

5.2.5 Values Correlate with Privacy Preferences Similarly

Lastly, we looked at the distances between how *Value* and *Value App* scores correlated with privacy preferences.¹³² These distances tell us how similarly values correlated with privacy preferences. This was conducted to assess to what degree values correlate in a manner consistent with the TBHV on which the survey questions were based. Some correlated similarly with privacy preferences, but many of these results were not what we would expect based on Schwartz's work (Schwartz, 1992; Schwartz et al., 2012). To demonstrate this, some results of note are reported here. *Achievement's* correlations were dissimilar to all other values' correlations (all Euclidean distances $d > 0.48$), where we would have expected a more similar relationship with similar values, such as *Power* ("Self-Enhancement"). In addition, for OpenLitterMap, *Security App* and *Conformity App* correlated dissimilarly ($d = 0.64$) – indeed, like *Achievement*, *Conformity App* correlated dissimilarly with almost every value ($d > 0.52$) where we would have expected a closer relationship with values such as *Tradition* ("Conservation"). For Lose It!, *Hedonism App* also correlated dissimilarly to most other *Value App* scores ($d > 0.4$), including *Stimulation App* ($d = 0.73$), where similar correlations would be expected (both "Openness to Change," Figure 2-2). While there were some instances that conformed with Schwartz's theory – for example, *Conformity* and *Security* showed close similarity with each other ($d = 0.068$) as we would expect based on their proximity on the Schwartz quasi-circle (both fall on the "Conservation" dimension) – the prevalence of non-conforming results suggests possible method limitations, explored in the next session.

5.2.6 Survey Conclusions: Theoretical and Methodological Implications

Considering the app-specific nature of the correlations between values and privacy preferences, we conclude that value-privacy preference relationships are app-specific. While all correlations were weak ($\rho < 0.5$), this could be because the cumulative group of privacy preferences relate to certain values, rather than individual privacy preferences. Such an interpretation also makes sense when considering previous personalized privacy assistant investigations into privacy preferences for profile building (Liu et al., 2016), which found the app category as one of the more significant indicators of allowing certain data collection practices.¹³³ In addition, while values are trans-situational (e.g., one; hierarchy of values is preserved), which values are relevant is context-dependent in the TBHV (Schwartz et al., 2012). Here, we are relating the two – privacy preferences and

¹³² Graphs available at: <https://zenodo.org/record/8208858>

¹³³ Liu et al. (2016), pg. 31: "We fitted the users' settings data to a random effect logistic regression model grouped on users' allow/deny decisions on app permissions. The independent variables include major features that could be obtained in our [Android user] dataset such as user demographics and app category. [...] App category and the type of permission are significant predictors for an individual's allow or deny decision."

How Do We Value Data Privacy?

values – in pursuit of a value-centered approach, and the observed context-dependency is in according with these works.

These results, however, also reveal methods limitations, and support the need to explore the value-privacy-app relationship by other means. In addition, the relatively weak correlations could indicate that the 10 values from the SVSS survey are not specific enough to tease out the relationship between values and privacy preferences. We also received a few emails from participants expressing confusion around the value questions on the survey, focused on understanding the value's definitions or struggling to comprehend their applicability to an app. While we presented preliminary interpretation of individual correlations here, these links are tentative and possibly due more to inconsistencies in participants' understandings rather than meaningful connections.

In future work, one could consider using the more comprehensive tool from the TBHV, the full SVS survey, despite concerns around the time it takes for participants to fill out. However, our results seem to suggest that switching to the longer SVS within the TBHV would not solve our issues. Accounting for context-specificity and presenting values in a manner that is understandable across cultures are known challenges of the SVS and the TBHV. High variability due to different interpretations of values, which can be heavily framed by one's culture, is also an ongoing challenge in cross-cultural psychology and the study of values (Karadag et al., 2018). The high variability we saw here in the dataset, the corresponding weak correlations, and the emails we received support this critique, at least in the case of measuring values in privacy or smartphone settings. In addition, it is also a central tenet of the TBHV that the order in which we prioritize our values is trans-situational, even if certain values are not applicable in certain contexts (Schwartz, 1992; Schwartz et al., 2012). The relative importance of applicable values is also viewed as critical for guiding one's actions. While different *Value App* scores between apps are in alignment with Schwartz's value theory, the correlation differences between Lose It! and OpenLitterMap seem to suggest hierarchies are *not* conserved. This result supports previous critiques of trans-situational values (de Wet et al., 2019), this time in the case of studying the value-privacy relationship.

In addition, we must be wary of reductionism by assuming that the rich and diverse landscape of human values can be *fully* captured in Schwartz's ten values, or the methodology derived from them. Evolutionary psychology forms the basis for Schwartz's theory (Schwartz, 1992). Applications of evolutionary theory to social phenomena, including the field of evolutionary psychology, has been heavily critiqued for trying to apply unifying principles to explain human behavior based on biological principles alone (Caporael & Brewer, 1991; Dupré, 2001). In addition, psychology as a field has been critiqued for its WEIRD (Western, European, Industrial, Rich, and Democratic) biases (Henrich et al., 2010), which are also relevant to this investigation of values. In particular, it has been critiqued for overgeneralizations outside the participant group or context (Henrich et al., 2010), experiments lacking an understanding of Non-WEIRD cultural assumptions and values (Baumard & Sperber, 2010), and bias introduced by WEIRD people's predominance in the field of psychology (Meadon & Spurrett, 2010). While the TBHV provided us with an efficient and validated means of initially exploring values in terms of quantitative data, the results must be further complemented with different methods.

In summary, we have seen here that values, privacy preferences, and app download decisions seem to be related in a primarily context-dependent manner. However, given our results suggesting limitations of the TBHV for our research purposes concerning smartphone privacy, an alternative investigation is needed to explore values in data privacy decisions outside TBHV-derived means. The aim of these investigations should be to

counteract these shortcomings by allowing for the complex, and sometimes messy, nature of human values to be captured under less methodical and theoretical constraints. Such investigation will be a critical means of further validating the context-specificity of the value-privacy preference-app relationship. It will also allow us to dig deeper into the individual particularities – motivations, understandings, as well as context – of these relationships.

In the next section, I present a richer, less theoretically constrained semi-structured interview analysis of *how we value privacy* that starts to accomplish these goals. The interviews also provided additional insights into how participants experienced the survey, which is also discussed in the context of possible limitations of Schwartz model and the 4DT model of value-centered privacy decisions.

Section 5.3 Semi-Structured Interviews

5.3.1 Section Overview

Semi-structured interviews were conducted with 18 participants to complement quantitative data from Phases I and II. Here, we present the results concerning **RQ1**: *how do we value data privacy?* (Table 4-1). We also discuss insights gleaned from the participants pertaining to the identified limitations of the survey data.

Firstly, concerning the survey, participants had an overall positive experience – finding the value questions challenging because it caused them to reflect on their values. However, some participants struggled with understanding value definitions in the survey and their applicability to smartphone apps, and commented on the lack of discrepancy between how they *wish* to act, and how they *really* act. The lack of clarity could have contributed to the high variability and weak correlations we saw in the survey data, solidifying our view that the TBHV and SVSS are not best suited for investigating values in privacy decision-making. Taken together, this feedback suggests that the results obtained from the survey are still relevant for understanding values and data privacy. However, they must be further compared and contextualized using other methodologies.

Secondly, we explored participant values to answer **RQ1**. While values involved with more explicit privacy decisions, such as online cookie consent notices or smartphone permissions requests, were interrogated, we were particularly interested in the values involved in smartphone app choice as implicit data privacy decisions of interest. Data was analyzed using a process of Reflexive Thematic Analysis (Reflexive TA), where values (themes) and sub-values (sub-themes) were *loosely* informed by the 4DT understanding of values. Based in part on this understanding, values and tensions between values were interrogated, coded, and grouped based on commitment statements¹³⁴ and affective responses.¹³⁵ Six values and fourteen value tensions were identified using this method, with their nuances described in this chapter. Participants strategies for resolving value tensions are also described.

Taken together, we conclude that these value observations suggest a complex, context-dependent motivational relationship between values and privacy decisions. Some values, such as *Control* (especially *power and choice*) and *Use* (*utility and function*), spanned contexts and were quite prevalent in the data. Besides these, *Community* (*connection*), and *Safety* (*security*), were also quite prevalent. In addition, we conclude that

¹³⁴ As a reminder, 4DT understands values as clusters of similarly oriented commitments on how to be and act in the world.

¹³⁵ Code list with references available at: <https://zenodo.org/record/8208858>

How Do We Value Data Privacy?

the 4DT value-centered understanding for value involvement in privacy decision-making does seem to capture the role of individual values in privacy decision making with, for example, the emphasis on the value *Control* supporting an understanding grounded in self-governance. Participants also expressed tension resolution strategies that could be improved to further support their ability to act according to their values. The line between value tensions that can be resolved and *double binds* was not clear-cut in the data, but rather, a spectrum depending on the participant's values and their ability to act. However, the concept of *double binds* was still helpful to identify areas where users were *most* unable to act in full accordance with their values – that is, were most “*double bind-like*.” In particular, we found that the presence of structural factors – such as social media monopolies, the attention economy, and surveillance capitalism – were defining features of these more extreme cases. Other potentially relevant interview observations, such as participant concerns on data-based business models and misunderstandings about privacy, are also touched upon in relation to how they relate with value-centered privacy decision-making.¹³⁶

a Participant Demographics

In total, 46 Mock App Store participants agreed to be recontacted for an interview. After being contacted via email, 20 agreed to be interviewed, and 18 interviews were conducted based on both saturation and informational power considerations.¹³⁷ Participants' stated gender was evenly split between women (9) and men (9) (Appendix VI).¹³⁸ Nearly all were adults (ages 25-64), with only one young adult participant. The majority also had (or were in the process of obtaining) a doctoral or master's degree (15) and 3 a bachelor's degree. Nationalities were grouped by continent, the majority with European nationalities (9), followed by Asia (6), and all others (3). Fluent and native English speakers were evenly split, and all participants currently owned a smartphone.

5.3.2 Online Survey Feedback Supports Previously Identified Survey Limitations

To begin, we will look at interview questions that pertained to how the survey data was perceived. The results from the online survey (Section 5.2) suggested possible limitations of applying the Short Schwartz Value Survey (SSVS) to explore values relevant to smartphone privacy. We therefore took the interviews as an opportunity to ask for additional feedback on the online survey. Roughly half (8) of the interview participants also took part in the survey. While feedback suggested some participants struggled with survey understandability, the feedback varied considerably between participants.

¹³⁶ When reading this chapter, there will be statements here that the well-versed privacy scholar may consider problematic misunderstandings about privacy – e.g., the “I got nothing to hide” argument (Solove, 2007). The goal here is not to call out misconceptions, but to note value-laden motivations for acting. However, I will briefly touch upon some common misunderstandings about privacy that appeared in the dataset and their relevance to the individual value-centered privacy approach in this chapter.

¹³⁷ See Section 4.4 for a greater discussion of saturation and informational power.

¹³⁸ Unfortunately, one participant completed the survey twice, resulting in 19 (instead of 18) responses. The pre-interview survey (Appendix V) was anonymous and not linked to an email or person, and therefore the duplicate could not be identified. We think that the person was 25-64, European, with a higher education degree based on the most likely participant pool. We therefore removed one from each of these demographics.

a Positive Survey Feedback

Many of these participants had a positive experience with the survey. Five reported that they found the survey questions clear and understandable. P14 further reported that they felt having “a sliding scale” for ranking values over a binary scale was helpful, and P01 said that they felt the number they gave for value was “a good representation” of how they feel. P01 further reported finding the links between values and apps clear, especially between *benevolence and universalism* and OpenLitterMap. Three reported that the value survey questions were somewhat difficult but in a way they found overall positive, because it caused them to reflect deeper or sometimes for the first time on how their values affect their privacy preferences and smartphone app choices (Q31, Appendix IX).

b Possible Limitations

Participant feedback that calls into question the reliability of the survey was around complex terminology and lack of distinction between values they feel they have and how they *really act*. Three reported finding the value questions difficult to answer because the terms used - such as “social power” – were not everyday terms and their meaning unclear to them. This contributed to further confusion regarding the app-specific value questions, with three reporting that it was not clear how the values related to the app. P07 found the value questions difficult because they were not sure whether the questions accurately captured their opinions and views on privacy. P13 did not realize that the privacy preferences were for a specific app, not apps in general. Two also reported that their values change based on context and found it difficult to answer the general value questions. In addition, P07 felt their answers to the value questions did not reflect their values because there is a discrepancy between how they act and how they feel, or their “ideal” (Q32, Appendix IX). The survey intended to capture what is “ideal,” even if not fulfilled now, to identify points where values were not being acted upon. These misunderstandings about the survey, taken with the feedback we received regarding the survey described earlier in this chapter, could have affected the results, and explain the highly variable and weakly-correlated value-privacy preference relationships observed in the survey results. These statements further support the limitations of using the Schwartz method to tease out the role values in privacy decision-making.

c Conclusion: Limitations and Relevance of Survey Results

Based on these results, survey participants had an overall positive experience – finding the value questions challenging because it caused them to reflect on their values and overall clear. However, some participants struggled to understand what the value questions were asking them – whether it be from value understandability, how they really act vs. how they wish to, or confusion regarding their applicability to apps. The lack of clarity could have contributed to the high variability and weak correlations we saw in the survey data, solidifying our view that the TBHV and SVSS are not best suited for investigating values in privacy decision-making. In addition, this suggests that the survey results need to be further compared and contextualized using other methodologies. Critically, though, the fact that many of the participants interviewed had a positive experience with the survey suggests that the results obtained are still relevant for understanding values and data privacy. While we may not be able to draw conclusive conclusions for each observed correlation, we can reasonably interpret the results of the survey to conclude that value-privacy preference relationships are highly context (app) dependent.

5.3.3 Six Value Themes: Complex, Context-Dependent Value-Privacy Relationship Grounded in *Control*

a Value Overview

To explore what and how values are involved in data privacy, semi-structured interviews were conducted. From these, six values and fifteen sub-values were identified as relevant to data privacy decisions and app choices (Table 5-2). While values (themes) could be constructed using a 4DT-undoing of values and Reflexive TA,¹³⁹ each individual participant's understandings and value-laden motivations for making certain decisions were quite varied. Relevant values also varied depending on context, although some of the more prevalent values, such as *Control* and *Use*, largely spanned contexts. Value theme nuances, individual differences, and similarities are presented and discussed here. Other potentially relevant interview observations, such as participant concerns on data-based business models (surveillance capitalism, the attention economy, and social media/app monopolies) are also briefly discussed.

We conclude that these value observations suggest a complex, context-dependent motivational relationship between values. Some values were more prevalent than others, with some of these values such as *Control* and *Use* spanning contexts. We further conclude that the relationship between values and data privacy decisions can be reasonably captured by a 4DT lens. We observed variations of expected phenomena that undermine autonomous, value-centered choice, such as the *weakness of will*, the “apathetic user,” the *inertia bias*, dark patterns, and a lack of privacy controls (Chapter 3, Tables 3-1 and 3-2).

(1) *General Observations*

Before diving into the results for individual values and tensions, there were some notable, overarching observations regarding how participants saw their overall relationship between values, privacy decisions, and apps.

Firstly, participants explicitly noted inconsistencies between their values and their privacy preferences or app choices. This supports reconceptualizing non-value affirming privacy decisions as value inconsistencies as described in the 4DT approach to data privacy (Figure 2-1 and Table 3-1). In particular, statements such as these suggest that users may *intend* to act in a value-consistent manner but fail to do so. This could suggest either *weakness of will* or presence of (inappropriate) nudges, all resulting in a failure of *external self-realization* and *self-unification*, over *akrasia* (forming intentions against one's personal identity and values).

Secondly, another participant (P09) stated that security is not the only value involved in their privacy choices, describing not harming others and utility as other critical considerations (Q1, Appendix IX). This further supports an understanding of privacy decision-making as involving multiple personal values, captured in the “digital home” analogy in Chapter 2, where our apps and the data we share with them are like furniture or people we allow into our home.¹⁴⁰ The “digital home” understanding is supported by statements concerning multiple values.

¹³⁹ See Section 4.4 for a detailed explanation of the coding method.

¹⁴⁰ Explained in greater depth in Section 2.4.

How Do We Value Data Privacy?

Values	Core Concept	Sub-values	Example:
1. <i>Community</i> 	"Other-facing," social and world connection	1. <i>accessibility</i> 2. <i>authenticity</i> 3. <i>benevolence and universalism</i> 4. <i>connection</i>	P06: I am in contact with [friends] on WhatsApp , which I very much appreciate [...] we keep very, like very connected even though we haven't seen each other in two years [...]
2. <i>Control</i> 	Freedom to make and the power that restricts choice	1. <i>dignity and respect</i> 2. <i>power and choice</i> 3. <i>tolerance</i>	P15: It's like a war , right. I mean, who's going to win? I mean, I'm sure [the companies] are going to win one way or another way, but I'm not going to give them an easy win .
3. <i>Growth</i> 	Improving oneself and meeting goals	1. <i>learning and staying informed</i> 2. <i>self-improvement</i>	P22: I think [the app] shows me probably things that I value , like how much exercise [I have had], it's kind of a nice reminder [...] And I guess this feedback is important in helping me improve [and] live a better lifestyle.
4. <i>Pleasure</i> 	Enjoyment, hedonism		P07: There's been like, a couple of games that I've downloaded to like use on the airplane and things like that. And [...] they're just not fun , so I'll delete them.
5. <i>Use</i> 	Convenience and function	1. <i>time and convenience</i> 2. <i>utility and function</i>	P06: I choose [apps] for functionality [...] I have one game that I regularly play. And most of the others they like, use some digital linkage with my gym [and] my university. So [...] my phone is very functional .
6. <i>Safety</i> 	Safe and secure in personal life and world at large	1. <i>non-maleficence</i> 2. <i>security</i> 3. <i>trust and trustworthiness</i>	P26: As I go like, "only the absolute necessary [cookies] " [...] simply because people can track you [...] you're not doing anything illegal [...] but the fact is that identity fraud is a big problem [...] people's identities have been stolen .

Table 5-2: Summary of value themes and sub-values from semi-structured interviews

Thirdly, some participants noted that what values a company represents are important when deciding whether to download an app or engage with a service. P01, for example, discussed how they “like the values [Firefox] the company represents.” Another app they “use a lot is Signal,” because they are “better privacy protecting.” P01’s comments in particular are interesting because Nurwidyanoro et al. (2002) found that app value statements – for example, stressing the value of privacy for Signal and Firefox Focus – seemed to influence the values they identified when exploring relevant values for apps on GitHub discussion forms. This suggests that the values one holds and their synergy with company value statements can be influential when looking to choose value-consistent apps,

How Do We Value Data Privacy?

further suggesting a role of personal values in choosing an app and more general privacy decisions.

Lastly were perceived (by the interviewee or interviewer) demographic influences on the relationship between values and privacy preferences. Five participants felt that one's age – being younger (P2, P6, P12, P26) or from a “different generation” (P26) – had a profound effect on their current and past privacy preferences, values, or both, as well as on those around them. Those who worked with personal information in their occupation also frequently expressed the central role privacy and data protection plays in their working life, and how this has shaped their own perceptions of privacy. This was especially true for those in research, where participants expressed a high degree of responsibility for data protection and privacy in order to respect individual data subjects (Q2, Appendix IX). In addition, participants with tech-related research experience tended to discuss what the “user” would like when asked about values and privacy, rather than what they themselves would like. However, working in technology did have a profound impact on one of the participants (P15). They described a “cynical” attitude towards apps that want access to their data, limiting apps on their phone to those needed for their work and heavily scrutinizing apps that request “questionable” permissions.

P15: “I overthink when it comes to [downloading apps] maybe because I worked in the background of it.”

Statements such as these suggest a greater role of one's personal attributes and professional life in shaping their *self-defining* attitudes (beliefs, goals, values, and commitments; Figure 2-1 and Table 3-1) that make up their personal identity. While I will not be teasing apart how we form the values we have,¹⁴¹ it is still interesting to note due to the prevalence of this theme in the data; perhaps this could be an avenue of future research.

These general considerations aside, the value-privacy relationship is best interrogated on a value-by-value basis to capture the myriad of understandings and the nuanced, largely individualized motivational relationship between one's values and privacy decisions. We'll start by looking at the most prevalent six values - firstly, *Community*.

b Community

The value of *Community* was focused on social and world connection; it was “other facing.” There were five sub-values: *accessibility*, *authenticity*, *benevolence* and *universalism*, *conformity*, and *connection*.¹⁴² While this general value was quite prevalent, some of the sub-values were observed with only one or a handful of participants. Even more prevalent sub-values, such as *connection*, could be understood in different manners. In addition, *Community* values appeared both in the context of data sharing and app choice, and were largely, but not exclusively, related to apps and services that had a social component. This suggests that *Community's* relevance was quite context-dependent, and the applicability of its sub-values could be quite varied by individual.

¹⁴¹ See footnote 40.

¹⁴² To distinguish between the two, values are presented capitalized and italicized (*Values*), and sub-values lowercase and italicized (*sub-value*).

How Do We Value Data Privacy?

(1) *Accessibility*

Accessibility focused on data and technological implications on the most vulnerable. This sub-value, however, was only relevant to one participant, P26. They described, for example, how tech should be used to make government services more accessible for vulnerable communities but currently is not (Q3, Appendix IX). P26 also expressed concern that long terms and conditions when consenting to data privacy sharing are especially challenging for those who are more vulnerable. This sub-value was held in high esteem by P26 but was not held by others – suggesting that *accessibility's* relevance is highly individualized.

(2) *Authenticity*

Authenticity focused on the meaningfulness and value of our online and offline interactions. *Authenticity* was expressed as a positive orientation to non-tech mediated engagement with the world - a tension between the negatively-valenced *digital world* and the positively-valenced *real world*. It is more meaningful, with the online world, in contrast, described as just one “large abstraction” (P08). One participant also noted how this abstraction affected their data privacy decisions online, with the effects and risks of their data sharing so far removed from everyday life.¹⁴³ Participants also did not like app notifications or video games distracting them or pulling them out of their non-digital life, with one participant expressing taking regular periods of disconnect to get back in touch with their real life. In contrast to the others, however, it is notable that one participant expressed that apps *help* them engage with the authentic *real world* by using Google Maps “for exploring” (P13).

*P13: “[Google Maps] is the one that **allows you to discover places**, because it has like little signs and things to see, restaurants, or it calculates the how long it's going to take for you to cycle from this point to this point and choose [...] appropriate path [...]”*
[Emphasis added]

These results suggest that *authenticity* may be differently understood based on the individual, causing them to see the digital world in conflict or as a means to better engage with the real world. This perception further motivates them to engage with and share data with applications or refrain from engagements depending on how they view *authenticity*.

(3) *Benevolence and Universalism*

Benevolence and universalism centered on choosing actions that benefited others and the world. When it came to benefiting the world, participants noted how profiling and data sharing can help society as a whole - for example, one participant felt that notifying people of nearby convicted sex offenders was beneficial for the community. Supporting the developer for some well-made apps was also viewed as promoting good in the world. One participant also noted how apps such as OpenLitterMap can help us get over our initial inertia to help others and our world by making it easy to help, creating a “low barrier” and thereby promoting *benevolence and universalism* (P01). In contrast to this, and closely related to *authenticity*, was the idea that disengaging with apps could also be seen as valuable for engaging beneficially with nature and the world. This is especially interesting,

¹⁴³ This is more relevant to another sub-value, *power and choice*, discussed in (4) *Control*.

How Do We Value Data Privacy?

given that we observed greater relevance of *benevolence and universalism* for OpenLitterMap over LoseIt! in the Phase I survey data. It suggests that this sub-value of *Community* can be context dependent. It also suggests that one's individual understanding of how *benevolence and universalism* is promoted or hindered by one's data privacy decisions affects one's choices.

When it came to doing good for others, the emphasis was that data collected by an app or online *should be used to benefit* the user in some way, whether it be mindfulness apps like Calm and Headspace, period tracker apps, or providing relevant ads for products of interest (e.g., books from Amazon; P01: "And they're also good recommendations. They also truly recommend books that I enjoy."). How beneficial a targeted ad was also depended on what it was promoting. For example, ads for an "educational institution" (P30) or helping small business owners were considered beneficial for society, while Facebook promoting "political ads" were "a bit perverse" (P01).¹⁴⁴ The view that apps and data should be used to benefit the user and promote overall good in the world was a fairly consistent understanding of *benevolence and universalism* across participants, but there was no general consensus on what the "good" being promoted is. Not everyone was as positive as P01 about targeted ads, even when recommending relevant products (discussed more in relation to *power and choice* below). This suggests additional individual nuance to *benevolence and universalism* based on one's understanding of what uses of data are beneficial.

(4) *Conformity*

Conformity centered around *fitting in* with what others were doing when it came to apps or sharing data. Central to this was the feeling that times are changing when it comes to downloading apps, and that we must change with them. To not do so would make you a social pariah.

P08: "I mean, **[Google] gives us an amazing access to information**, you know, I don't think somebody [...] could reasonably go about the world now **without access to that information**, without kind of **maintaining themselves as a stone age type person**, you know? [Pause]. I can't. I can't think of anybody I know who doesn't use Google."
[Emphasis added]

The sub-value spanned multiple contexts and had similar *overall* effects on app and data sharing behavior when present. Participants would delete apps if an insufficient number of peers were using them, and blogs, influencers, and friend recommendations had a strong effect on which apps participants downloaded or shared data with. However, some participants gave more weight to different groups when valuing *conformity* – for example, alignment with blogs and influencers were not always considered. This suggests that who one is doing the conforming to – and therefore the *specific* effect of conformity on their behavior – was more individualized.

¹⁴⁴ Likely a reference to the 2018 Cambridge Analytica scandal, touched upon briefly in Chapter 2. For more details, see a collection of reporting by *The Guardian* at: <https://www.theguardian.com/news/series/cambridge-analytica-files>

(5) *Connection*

Closely related to *conformity* is *connection*, which focuses around forming and maintaining interpersonal relationships. Particularly prevalent was the use of WhatsApp for connecting with others, especially for those who lived away from loved ones.

P06: “I am in contact with [friends] on WhatsApp, which I very much appreciate [...] we keep very, like very connected even though we haven't seen each other in two years [...].”

Social media services and apps were also mentioned frequently for their connective value, including connecting with others with similar interests or hobbies. This was viewed as beneficial for both personal and professional relationships and was especially valuable during the COVID-19 pandemic. While largely associated with social media and messaging services, the value of sharing and supporting each other to meet health goals using health apps was also expressed, as were apps for translating and communicating with others in another language. This suggests that *connection* is quite specific to social apps but can also be associated with other apps that serve a particular social function based on their individuals' needs, such as communication in another language.

c Control

Control focused on the freedom to make choices and the power that restricts choice when it comes to our digital interactions. This includes online interactions, apps, and data privacy choices. There were three sub-values: *dignity and respect*; *power and choice*; and *tolerance*. It was one of the most prevalent values and its relevance spanned contexts and individuals, suggesting that *Control* is a highly relevant value to privacy decision-making. Its prevalence, and discussions around especially *power and control*, also support the 4DT approach to value privacy. There are instances of expected phenomena, such as the “apathetic user,” where a participant becomes unwilling to take on commitments pertaining to their data privacy. However, discussions of *Control* provide some additional insights into how these phenomena come to be, such as the “apathetic users” arising not just from notice fatigue, but from a sense of futility that results from existing power structures.

(1) *Dignity and Respect*

Dignity and respect encompassed the belief that humans are not data points; they have inherent dignity as human beings and are deserving of respect and consideration. This was especially relevant when it came to data sharing and its effects. For participants, there should be limits to what (sensitive data) can be tracked. Companies should also use data how they say they will and not be “sneaky” when collecting it (P15).

*P15 “[...] It's [an] ethos [...] I know that [...] I'm getting this product for free. So basically, you know, that saying that “if you are getting it for free [...] [then] you are the product.” So, I understand these **companies have to make money**, but [...] **they can't be sneaky**. That's not right. That's something that **I can't agree with at all**. If you want [my] data, make sure that [...] you are very clear about that.” [Emphasis added]*

In terms of targeted ads, a handful of participants felt disrespected by being profiled for targeted ads, proposing that ads should be less persistent and invasive. Like *authenticity* (*Community*), it was expressed that data should also not be used to impact one's offline

How Do We Value Data Privacy?

life, such as to influence political beliefs. However, unlike *authenticity*, this was considered unacceptable out of respect for that person as an individual, rather than being inauthentic.

Statements pertaining to *dignity and respect* demonstrate that users see themselves and others *as individual agents deserving of respect*, the normative basis of the 4DT lens utilized to understand value-centered privacy decisions (Section 2.4). Their statements further emphasize how their *dignity and respect* as agents is being violated when it comes to data sharing, further bringing home the need for our value-centered understanding of 4DT.

(2) *Power and Choice*

Power and choice was one of the most prevalent sub-values throughout the interviews, with its relevance spanning contexts and individuals. It is also most relevant when considering the 4DT approach to value-centered privacy. *Power and choice* focused on one's agency over their data privacy and app decisions as well as the power structures that may manipulate or coerce their choices. In this case, the power to manipulate was understood as an external other (someone) encouraging a choice against the user's intentions. This manifested in both app choice and data privacy decisions, with participants expressing varying degrees of perceived control and manipulation. They may have felt a duty to self-govern, a sense that is harmonizable with 4DT's emphasis on taking on commitments and engaging in the process that constitutes being an autonomous agent (Figure 2-1, Chapter 2). However, while they had a sense of responsibility for their data privacy choices, many were ultimately unable to fully act (*self-realize* and *self-unify*), due to existing structures that restrict and manipulate their actions. As we will see, some even reach the point of apathy – the “apathetic user” – not just from notice fatigue, but from the futility of acting in our existing privacy environment.

In terms of data privacy decisions, many participants discussed the degree of control they felt they had over their choices to give away data. Despite some participants' intentions, control was perceived as low – terms and conditions are impossible to read; the implications and risks of sharing data too abstract (especially when AI was involved); and privacy notices “specifically designed” to get us to consent, with “forms and designs [...] just too clever for individuals to [...] overcome” (P08) (Q4, Appendix IX). For a few, apps and online services were also viewed as being made to be addictive in order to get us to share more and more data with them. One participant even mentioned how, despite their best efforts to be mindful with data sharing online and on their smartphone, they were concerned that companies have ways of getting at their “subconscious” (P22). These instances constitute a mixture of (inappropriate) nudges and notice fatigue, failures of *self-realization* or *self-unification* resulting from crafty “dark patterns” and long, fatiguing terms and conditions. Two participants also expressed the desire for more buttons or mechanisms to control their data, as they were unable to act in the way they intended. This “lack of relevant controls” constitutes a failure of *self-realization* and *self-unification*, detailed by a 4DT lens to data privacy as occurring when acting according to one's intended, value-aligned privacy preferences is not an option.

It is not surprising, then, that many participants viewed trying to control their data privacy as futile. For many, they were in too deep already with how much data they had already shared into the world, and no matter what they try to do, companies have ways of getting at their data. Participants described their attempts to keep their data safe, but to no avail, such as using Mozilla Firefox but still finding Facebook logs in their browser history (Q5, Appendix IX). Participants also felt that companies, especially Meta/Facebook and

How Do We Value Data Privacy?

Google, use their data in ways they did not consent to. For some, the inevitability that their data will be collected and used outside their control leads to a sense of apathy.¹⁴⁵ No matter what, companies already have their data and the means of getting more.

*P02: “Sometimes I just [click] whatever [button] comes, [the] closest button to dismiss [the privacy notice]. [...] Google's probably running half our browsers and tracking us through that anyway. **It feels very futile to me.**” [Emphasis added]*

In these cases, these participants decided to *no longer care*¹⁴⁶ – take on commitments, deliberate, and act – according to their values. This is a failure of *self-constitution*, originally discussed in the context of third-degree notice fatigue. In this case, while this failure was brought on for some by the sheer unreadability of long terms and conditions, many came to be “apathetic users” due to the *power* of tech companies to get and do whatever they wish with their data, no matter what. This is an interesting finding, and not inherently contradictory to 4DT. It is important to note, then, that failures of this kind can result from power structures, and not just notice fatigue.

In contrast to those who became “apathetic users,” a select few participants viewed their futility in the face of powerful companies as a call-to-arms – even to fight in a losing battle.

*P15: “[Managing my data privacy is] kind of **like a war**, right? [...] who’s going to win? [...] I’m sure [the companies] are going to win one way or another way, but **I’m not going to give them an easy win.**” [Emphasis added]*

In these cases, such as P15, the user refuses to become apathetic and continues to take on and hold their commitments. However, their ability to realize what they intend is highly restricted – they are, in summary, *self-constituting* but not (externally) *self-realizing* and *self-unifying*. This, too, is not inherently in conflict with 4DT – in fact, it can be understood as a lack of relevant controls or course of action. However, these results suggest the centrality of power structures in restricting that choice.

This ties into another common concern among participants: how data that is collected online or on their smartphones could be *used* to shape their choices in the future, their “future tense” and their fight to preserve it under surveillance capitalism (Zuboff, 2019). In these cases, *power and choice* was less focused on the power and choice over the privacy decision itself, but how the data shared could be used to lessen their power and choice in other aspects of their lives. This focused mainly on targeted advertising and the power that comes with data centralization. To the former, many participants expressed that they felt like they were being manipulated by targeted ads, which they had little control over and that utilized personal data in ways that increasingly influenced their lives online and offline. Advertising AIs are “a little bit too smart,” with “a mind of their own,” putting all these “prompts in [their] life” to “sway [their] decision [...] [to] what [these] shops want [them] to do” (P22) (Q6, Appendix IX). One participant felt like these ads intentionally misrepresent the product. Besides manipulation by targeted ads, two participants also mentioned that they were wary of ever sharing their phone number

¹⁴⁵ What I am calling the “apathetic user” has also been described as “Privacy Cynicism.” While I am aiming to understand how out apathy hinders value-centered privacy decisions, it is well-documented in the privacy literature that privacy cynicism can influence one’s privacy decision-making behavior. For example, see van Ooijen (2022).

¹⁴⁶ This is key when distinguishing an “apathetic user” from someone in a state of “digital resignation” (Draper & Turow, 2019) . See footnote 63 for more details.

How Do We Value Data Privacy?

because of how much easier it is to be manipulated on the phone by a particularly crafty communicator. Participants also frequently expressed concern around data centralization, especially the amount of power it grants those who hold that data, such as Google. Some participants mentioned the power this grants Google to lock them out of their account or use this centralized data to potentially prevent them from getting a job.

*P01: “[Centralized data] gives less power to each individual agent. And **if tomorrow, Google turns evil, I will be in big doo-doo.** I will be in trouble. Because **they have a lot of information about me.** But if tomorrow the company I store my passwords with goes evil, okay, it will be problematic [...] but I will still have access to the to the company that has that stores, my two factor authentication keys, I will still have access to my email provider, I still have access to my maps provider, **but if they're all the same company, then its centralized power is always scary.** [...] it doesn't need to be political [...] what if Google releases [...] [a] business tomorrow [...] **an app for businesses where they can look up any person on the planet,** and they can obtain a file [...] before hiring somebody. They're going to look up the name of that person and see what they have posted online. And if when company has all that information, then [...] **it's already scary, it doesn't necessarily [need to be] political. [It could] just be whether or not 10 years ago, you posted the picture smoking,** and that company that's going to hire you is now **going to discriminate [against] you or not or going to hire you [...]** I feel that in a more decentralized world, that's going to be hard [...] to accomplish, because [...] [no] single company is going to have all the data.” [Emphasis added]*

Other participants mentioned the *power* of centralized data to predict their behavior and profile them. Using this data, companies have “formed such an effectively predictive algorithm of [our] behaviors, that they know what [we] want before [we] know [we] want it or they know what [we’re] going to discuss before [we] think of discussing it” (P09) (Q7, Appendix IX). In line with concerns over prediction power expressed by P09, another participant expressed frustration with how plane tickets seem to change based on their search history and how this perceived profiling had pushed them to use Incognito mode on Google Chrome. Other participants also expressed concern on how data centralized in social media monopolies can be used to (emotionally) manipulate us. Furthering this power dynamic was an overall sense of lack of transparency¹⁴⁷ around how these social media AI work. One participant aptly captured these elements by comparing the social media experience to playing poker.

*P08: “I guess I did go “all in” to social media at one stage [...] But I felt like... [pause]... I felt like I used to play a lot of poker. And **when you[re] sitting' at a poker [table], [...]** **you're either the one kind of controlling the ideas or the situation or you're the one being controlled,** [...] [or] other people are dictating your emotions at the table [...] and I **felt like, online, I wasn't having control of my emotions.** I felt like “[...] what the hell made me angry there? Why [...] [is] my blood [pressure] rising when I'm just reading*

¹⁴⁷ While not directly related to this work, it is also interesting to note the concern P26 had around not knowing when they are talking to a bot or a human and the need for more transparency: “I was on a chat the other day was asking for [...] some information [...] I said, [the information I want is] in [home city]. And then suddenly, this, “Joe” comes back and says, “Oh, I was in [city] once. It's a very pretty city.” And I was like, “oh, yeah, Joe, when did you-?” And [it] didn't answer me. Because it's a bot. [...] And it's named Joe. You know, it has a human name. Like it's not clear that it's a bot. [...] people need to know that they're actually not talking to a human being.”

How Do We Value Data Privacy?

stories [...] there's an issue [...] with social media [...] it's quite manipulative. So yeah, I did [...] [come]off of them.” [Emphasis added]

This manifestation of *power and choice* – data used to manipulate us in ways outside our privacy choices – could be understood in traditional 4DT as hindering our *self-realization*, *self-unification*, and as inherently disrespectful of our human dignity (See Section 3.2). In the case of understanding value-centered choice in privacy decision-making, we can see in these results that users who care about *power and choice* are not always sharing data in a way consistent with this value. They give data away, that is then used to fuel AI that then affects their lives in some way. This would constitute a failure of *self-realization* and *self-unification* – again, tied up in the power dynamics of the systems such as surveillance capitalism.

Aside from expected control of data or how it could be used to restrict their choice, participants also discussed the degree of control they felt when it came to choosing and keeping apps on their phone. One participant expressed frustration that apps that come with Huawei Android phones cannot be removed, likely collecting data about them that they do not wish. Another participant felt that they “fall for the trap” (P29) when it comes to Instagram influencers encouraging them to download apps they would otherwise not have (Q27, Appendix IX). Interestingly, in this case, the participant felt that their value for say, health, may be used by the influencers to download an app they do not need (Q8, Appendix IX). Participants also frequently expressed frustration around mandated apps and online services for their work, such as Microsoft OneDrive (Q9, Appendix IX), although two participants notably expressed that they were not particularly bothered or upset about being asked to have certain apps and services for work, viewing it as a reasonable work expectation. There was also frequent discussion on the role of *inertia* in keeping and using apps. We previously met the *inertia bias* (a kind of nudge) when it came to privacy permission defaults and PPA profiles (Chapter 3, Table 3-2). We also see it here when it comes to keeping previously downloaded apps on our phones. Many participants have had the same apps for years that have already collected data on them. “It's gone too far,” and a “clean slate” is needed in order to “look at [apps and data privacy] a lot differently” (P29).

Regardless of online or on their phone, participants had developed strategies to fight for their “future tense” (Zuboff, 2019): try to reclaim control over their data, how their data is used, and the apps on their phone. From a 4DT lens, we can view this as methods of trying to reclaim their *self-realization* and *self-unification* by acting in accordance with their *power and choice* value – both *power* over the action of sharing data itself, as well as the *power and choice* implicit in the results of the data sharing (how the data is used). In this case, strategies were quite varied depending on the individual, and mostly focused around how to disrupt the influence of targeted ads or combatting social media manipulation. These included: using ad blockers, only accepting essential cookies, paying to remove ads from an app, doing their best to read the terms and conditions, carefully vetting apps, turning off ads on Android, “just swipe [the ad] away” without looking at it (P04), clearing cookies regularly, deleting apps that monetize sensitive data (e.g., period apps), consciously not downloading ads that have been advertised to them on social media, avoiding apps and services with unnecessarily obtuse “legal speak” (P26), deleting social media apps, only using social media on their web browsers, and having a separate phone for social media apps. Others actively “try to confuse the machine” when interacting with online targeted ads (P22) – although not fully clarified in the interviews, presumably this involves clicking on ads randomly or in manners that are not relevant to

How Do We Value Data Privacy?

them in order to feed the targeted advertising AI false data.¹⁴⁸ This is reminiscent of previous calls to *obfuscation* – the deliberate utilization of misleading information to combat data collection (Brunton & Nissenbaum, 2015).

Participants also expressed what regulatory or design choices could be made to improve their *power* to make choices about data sharing, data use, and apps. Participant views on tech companies, business models, and regulation heavily influenced how they felt best *power and choice* could be promoted. While nearly half of participants viewed data and their privacy associated with it as a tradable commodity for a good or service (P14: “you know, if you're using a free application, you're paying it for it in some way”), many participants were highly skeptical that: 1.) companies have their best interests at heart due to profit incentives;¹⁴⁹ 2.) were uncomfortable with data-based business models; and 3.) felt that some form of data privacy regulation should be in place. For these reasons, regulation was viewed as a means of reigning in data-based economies and companies that may use data in a manner that violates one’s *power and choice*. There should be a limit to what companies (and, for some, governments) collect, and certain kinds of data and uses should *not* be allowed. For example, P30 stated that “anything that could affect your personal freedom [or] your personal well-being, that should absolutely not be available,” and P02 felt that data should never be used to “sense that I'm sad” or “been drinking” in order “to start advertising me junk food.” Governments were also viewed as having the responsibility to educate citizens about data privacy and protect their data in the face of data monopolies and tech companies.¹⁵⁰ A few participants described existing laws such as the GDPR as a sometimes tedious but necessary mechanism to empower users and raise privacy awareness. This was especially true for those who had moved to the European Union from other countries and appreciated “hav[ing] the power” to disable cookies (P13; Q30 Appendix IX) – the regulation has increased their ability to *self-realize* and *self-unify* when it comes to their cookie choices. However, one participant felt that cookie consent notices could improve their ability to act as they intend to in their data privacy choices by being consistent, with the same cookies and definitions to uncheck.

While some viewed targeted ads in a more positive light than others, half of participants expressed that there should be mechanisms to have more control over targeted ads as a means of promoting *power and choice*. Some of these participants simultaneously noted how targeted ads can be useful but expressed that they “would prefer to just make [their] own decisions and choose how [they are] influenced” (P30). P09 had the innovative idea of a government-sponsored personal dataspace for citizens, a “best of both worlds” approach where one can decide which companies to allow access to their data while getting the benefits of targeted advertising (Q10, Appendix IX). In contrast, three participants expressed a desire for less or no targeted ads altogether, preferring “random” ads to preserve their choice. One participant also noted that education and awareness may help one manage targeted ads, noting that having some understanding of “psychology and philosophy and reading Shoshana Zuboff”¹⁵¹ (P08) may help them resist targeted ads more

¹⁴⁸ While not directly related to this work, it is also notable that one participant avoids answering the phone when it is not a number from their contacts to protect themselves from scams.

¹⁴⁹ For three participants, the business models of companies were fundamentally in tension with data privacy. One participant can understand why data is collected “from a capitalist point of view,” but does “not agree with them” (P22).

¹⁵⁰ Although they have the *responsibility* to ensure their citizens’ data privacy is reasonably managed, some participants were skeptical that they would *actually* act on it. For example, P26 felt that there is a “high level of privacy because it's government, but they do not seem to be interested in [...] [the] members of their state, you know. "It's okay for us, we have to have a very high thing, but we do not really care about anybody else.””

¹⁵¹ Likely a reference to Zuboff (2019).

How Do We Value Data Privacy?

than others, although they conceded that they see “how effective they are towards other people” who may not understand the dangers of targeted ads. It is also interesting that some participants had realizations about their data privacy during the interviews,¹⁵² including making new connections.

Another common view on how to promote *power and choice* was around transparency. For example, one participant expressed that companies could collect data but “they can't be sneaky” about it; they must be fully transparent (P15). Participants mentioned that what data is collected and how it will be used should be made clearer and more understandable, in “very simple, extremely simple” terms (P12). Participants also frequently acknowledged the current challenges of obtaining this in practice. In particular, two participants noted the challenges of AI transparency and explaining how our data is used by algorithms (e.g., social media feeds) as desirable but difficult to obtain. P09 even brought up the “transparency paradox” – the tension between understandability and completeness (Nissenbaum, 2011).¹⁵³ Despite this, participants did have some thoughts on how to present information in a way that would be more engaging and understandable for them, allowing them to better act according to their wishes. For example, two participants were in favor of more creative, interactive ways of presenting privacy information. For these participants, “it's not just [about] being readable, [it's about] being interactive, intuitive, and extremely simple” (P12). This could include “creative ways to help people understand, like graphics, infographics [...] rather than just [...] pages and pages of words” (P22).

It is notable that some participants found existing mechanism to exercise control over their data, targeted ads, and apps on their smartphones are at least somewhat helpful – that is, already allowing them to act in accordance with their valuing of *power and choice*. These participants felt that privacy notices online give them at least some power to say “no” to data collection. It came down, then, to a matter of personal responsibility – one’s responsibility to govern oneself to the best of their ability, even if aspects of their autonomy may be less than ideal (e.g., an insufficient number of privacy controls that limits their ability to act in full accordance with their *power and choice* value (*external self-realization* and *self-unification*). While these notices can be “annoying” (P10, P22), they were necessary to inform them of what is being collected and to exercise (some) control over their data. Terms and conditions were also viewed as long but necessary, and one participant even stated that they try to read them. Data collection, in these cases, is viewed as acceptable with this consent, and participants expressed that it is the individual’s responsibility to manage their data privacy.

*P10: “You can choose whether you are going to care about [privacy notices] or not. [...] When the pop up [appears], then you have the choice. And that's your choice. **You have to be responsible for that.**” [Emphasis added]*

Relatedly, targeted ads were not always viewed as manipulative. One participant stated that we still have some control over targeted ads that we can unsubscribe from (e.g., Amazon) – the power and choice over whether one gets the ads in the first place. In addition, the choice concerning whether to buy the product was still ultimately up to the individual.

*P06: “[You could say that] that [if] Amazon sends me very good advertisements, [...] [this] negatively impacts my financial situation, because, like, it induces me to spend money, but it doesn't. But **I still have some [...] control over it. Amazon doesn't go and***

¹⁵² Participants also had similar realizations during the Mock App Store exercise, described in Chapter 6.

¹⁵³ See footnote 20 for an explanation of the “transparency paradox.”

How Do We Value Data Privacy?

[...] just send me stuff. [...] [Or]” oh, yeah, we also have your bank details. So, we [...] [can] withdraw the money from your account [now]” [...] that would be a whole other step if they just started [...] preemptively doing that.” [Emphasis added]

Companies were still understood as doing their very best to sell products – but, like ads we encounter offline, we still have the choice to buy the product or not.

In summary, the *power and choice* sub-value of *Control* is fertile soil in which to explore the four dimensions of autonomy in privacy decision-making. The overall emphasis on *power and choice* (or lack of it) and orientations towards commitments and motivations, for both more traditional privacy decisions and app choice, are consistent with a 4DT understanding of value-centered choice. In addition, whether thinking we should manage our data privacy, even in a hopeless battle, or engaging with targeted ads, *power and choice* captures a sense of personal responsibility to govern oneself and act according to one’s values. This idea is harmonizable with 4DT that sees agents as needing to *self-constitute* – that is, engage in the process of taking on commitments to govern oneself (Figure 2-1, Table 3-1). However, participants also found “dark patterns” that nudge them to engage with an ad, share data, or download an app minimize *power and choice*. The companies are simply too big to protect your data from. This appears to lead, at times, to the “apathetic user” phenomenon – where the sense of futility means you throw up your hands and give your data away. It is interesting that this “apathetic user” phenomenon – understood using 4DT as associated with notice fatigue (Chapter 3, Table 3-1) – is associated with an overall sense of loss of agency under existing power structures. These power structures restrict an agent’s ability to *self-realize*, *self-unify*, or, in the case of the “apathetic user,” *self-constitute*, or a combination of the three depending on the context and extent of the existing power dynamics. Tackling these power structures, and thereby allowing all four dimensions of autonomy to be promoted, will likely need regulatory or broader interventions that allow users to exercise the degree of *power and choice* they hold when making data privacy decisions. Its prevalence also suggests that the *power and choice* sub-value of *Control* is a highly relevant value to privacy decision-making.

(3) *Tolerance*

The final sub-value of *Control* was *tolerance*, focused on accepting that others’ data sharing choices and technological interactions may be different from one’s own. This understanding was relatively universal among participants. Most participants found *tolerance* as a relevant value when it came to privacy decision-making – although this was not always the case. Those who valued *tolerance* felt that everyone should be able to share data according to their preferences. In addition, even if someone interacts with technology differently, participants felt they should tolerate others’ preferences. In contrast, the few participants who expressed concerns about *tolerance*’s relevance were more critical of others’ data privacy choices, having some understanding of what it is to manage one’s privacy properly. They felt that, for example, most people are too worried about their data privacy when the reality is that risks of sharing data are low, or that people can be hypocritical about their data privacy choices. While still valuing *tolerance* to a certain extent, these participants felt *tolerance* had its limits. This suggests that *tolerance*’s relevance is greatly related to an individuals’ opinions and beliefs around data privacy.

d Growth

The value *Growth* was focused on improving oneself and meeting one's goals. It had two sub-values: *learning and staying informed* and *self-improvement*. This value was highly context-dependent, constrained to apps and services related to education, news, social media, and health. This suggests that *Growth* is relevant to privacy-decision making in a context-dependent manner.

(1) *Learning and Staying Informed*

The sub-value *learning and staying informed* was focused on learning new skills and staying knowledgeable of world events. While there were some instances of FOMO¹⁵⁴ in the data, many participants instead found social media, news apps, and messaging services to be a vital and necessary place to get news in the modern world (Q11, Appendix IX). This was especially important for those who were living abroad, far away from family, because these apps allowed them to keep up to date on family milestones and country-of-origin politics. Apps were also opportunities to learn new skills, expressed by about half of participants. Duolingo was a frequent mention but podcasts, audiobooks, social media (such as Reddit, "a really useful community in terms of just gathering information." (P30)), and brain training apps were also mentioned. Taken together, *learning and staying informed* was quite relevant to social media, news, and education-related apps and services, with its relevance quite context-constrained.

(2) *Self-Improvement*

Self-improvement involves sharing one's data or engaging with an app or service in order to improve oneself. Nearly every participant interviewed (11/18) described using exercise, mindfulness, or other wellness apps to meet health goals, which often involve tracking health data.¹⁵⁵

*P22: "I think [fitness app] shows me probably **things that I value**, like how much exercise [I have had], it's kind of a nice reminder [...] And I guess this feedback is important in **helping me improve** [and] live a better lifestyle." [Emphasis added]*

Notably, *self-improvement* was nearly exclusively relevant in the context of health and wellbeing apps. As discussed with LoseIt! and the survey data (Section 5.2), the TBHV values *Achievement* ("success"), *Self-Direction* ("pursuing goals"), and *Stimulation* ("living a varied and challenging life") associations towards one's goals makes intuitive sense given the context – a health app to meet one's health goals. It is understandable as well that *self-improvement* would also be especially relevant in such a context. That said, a few participants did find *self-improvement* relevant in other contexts. For example, two participants mentioned the role of apps in helping them manage their spiritual wellbeing, such as apps with prayer reminders or a Bible app for daily reference. Taken together, these results suggest that *self-improvement* is very relevant to privacy decision-making in the context of health apps and is mostly irrelevant in other contexts.

¹⁵⁴ Fear Of Missing Out, e.g.: "I think there's that fear of [...] if you don't use [Instagram], then you can't possibly "be in the know"" (P07).

¹⁵⁵ This is perhaps not surprising given the content of the Mock App Store, which was used as the conversation starting point in the interviews.

e Pleasure

The value *Pleasure* focuses on the importance of enjoyment, hedonism, and engagement with apps and online services. It has no sub-values, nor did it relate too closely with the other values – its motivational effect on privacy decisions was quite distinct. In addition, *Pleasure* was mostly relevant to a specific context. It was frequently mentioned in the context of apps used for entertainment purposes, such as games, YouTube, social media, audiobooks/podcasts, and Netflix. *Pleasure* could also manifest itself through what the app *enables* them to do, such as pursuing hobbies for entertainment. It does, however, have a context-spanning aspect. It was important for participants that apps and services *themselves* be engaging and fun to use, even if their *purpose* was not entertainment (e.g., Duolingo). This was relevant to all apps and even the decision whether to engage with privacy notices. Apps with too many notifications were viewed by some participants as decreasing this enjoyment and, for three participants, apps that were not enjoyable would motivate them to remove the apps from their phone. In addition, finding cookie consent pop-ups that promoted feelings of frustration or displeasure was seen as a motivation for accepting all cookies and doing the quickest thing to get rid of them (P28: “swatting a fly”). In addition, while targeted ads were generally viewed negatively by participants for being manipulative (see *Control*), they were also viewed negatively for being simply annoying. Too many ads on apps were also a motivator for deleting apps or upgrading to the paid version of the app (Q12, Appendix IX). Taken together, it appears that *Pleasure* is quite motivationally distinct from other values. It is also primarily related to privacy decision-making in entertainment contexts, although considerations about an app, service, or notice’s overall level of engagement could also be of relevance.

f Safety

Safety focused on being safe and secure in one’s personal life and the world at large. It had three sub-values – *non-maleficence*, *security*, and *trust and trustworthiness*. While quite prevalent in the data, *Safety* sub-values tended to be more relevant in specific contexts, such as health apps or when sharing sensitive data. Understandings of *non-maleficence* were broadly consistent and context-constrained to social media and health apps. *Security* was primarily related to protecting one’s sensitive data from unintended and particularly harmful uses. In contrast, individual indicators of *trust and trustworthiness* were varied and resulted in different understandings of who is a trustworthy entity with whom to share their data.

(1) Non-Maleficence

Non-maleficence was oriented toward the idea that data collected as well as apps and services themselves should not be utilized to harm others or society. As we will see, understandings of *non-maleficence* were broadly consistent, with context-specific relevance to health and social media apps and services. However, harmful uses of targeted advertising spanned contexts.

Unlike *benevolence and universalism*, where understandings of what “good” should be promoted were quite varied, understandings of what constitutes harm were quite universal. For example, it was generally understood that if an app was being used to hurt others, such as the “original Facebook [...] rating women,” was “disgusting” (P09). This participant also stated that avoiding harmful applications like Facebook – that is, acting

How Do We Value Data Privacy?

according to their value of *non-maleficence* - was a critical consideration when deciding whether to engage with an app or service.¹⁵⁶

In addition, this value's relevance was also quite constrained to specific contexts, such as social media and health. Facebook/Meta was frequently viewed as violating *non-maleficence*, misusing data to cause harm to others and the world. This could be in the form of data-fueled politically-charged echo chambers or intentionally selling data for political purposes (Q13, Appendix IX).¹⁵⁷ To minimize this harm of social media and act according to their value of *non-maleficence*, one participant expressed that it was important to limit what you share on social media "because you don't necessarily know how they will use it." (P30). Participants also were concerned that social media could damage mental and emotional health. They viewed social media as a platform that allowed people to be meaner than they would be in real life. Participants also felt that the world looks much grimmer on social media than it probably is.

P07: "[...] maybe [social media] is not [...] what's going on in the world. Maybe the world like has always had bad news. But now we get it all the time. And it's everybody's opinion, and it's constant. [...] so maybe it just feels like the world's on fire when it's always been on fire." [Emphasis added]

Relatedly, health apps themselves could be used to harm. For example, one participant was worried that mental health apps do more harm than good, citing concern over access to mental health chatbots for those undergoing a mental health crisis. They also expressed skepticism at the "therapeutic value" of apps that want to "get [us] engaged" for profit (P02).

While *non-maleficence*'s relevance was particularly focused on social media and health settings, data being utilized for targeted ads was seen as a harm that may occur in many settings and contexts. One participant discussed how targeted ads can re-enforce problematic social norms and biases, possibly triggering emotional or mental distress.

*P30: "I like to think that [ads are] not intentionally used maliciously. But I know that sometimes you can [...] be targeted in a way that can be triggering for you. So like, one thing that I found kind of interesting is that I started getting ads for IVF.¹⁵⁸ And it's because somehow [...] they know that **I'm in my 30s. And I don't have kids. And they're like, "well, obviously, she needs to, like, have kids now. So, let's advertise IVF."** Like, this is so weird. It's not connected to anything I searched. Like, the only demographics they know about me is that I'm in my 30s. And they must have figured out that I'm childless. And I just think that it's really freaky, because like, they don't know my personal choices. And I'm deliberate about that." [Emphasis added]*

This participant also interestingly proposed that *why* the targeted ad is appearing can be a determinant of whether it is causing harm— is it because of a problematic assumption, or something more benign, such as location proximity to the business being advertised? While they did not believe that targeted ads themselves are purposely used to hurt others, they felt

¹⁵⁶ "[The original Facebook] is just disgusting to me. [...] you're using people and you're insulting and offending and humiliating people. So, I guess non-maleficence will be the first kind of value I'd endorse [when choosing an app or service] closely followed by utility/benefits. And then in third place, security [and] privacy" (P09). Also see: Q1, Appendix IX.

¹⁵⁷ Likely a reference to the Cambridge Analytica scandal. See footnote 144 for more details.

¹⁵⁸ *in vitro* fertilization

How Do We Value Data Privacy?

that answering the *why* question could be used by advertisers to prevent unintended harm and promote *non-maleficence*.

While understandings were mostly consistent, it is also important to note that other participants felt that the *lack* of observable harm to themselves or others meant it was acceptable to share their data. For these participants, sharing data online or on apps wasn't harmful because nothing bad had ever happened to them (Q14, Appendix IX). Contrary to others, restriction-less data was not viewed as inconsistent with *non-maleficence*.

Considering these results together, we see understandings of *non-maleficence* – that is, what is harmful, and the desire to avoid it – were mostly consistent and specific to social media and health app related-harms. However, some participants did not agree that limiting what you share was necessary to promote *non-maleficence* due to what they considered the lack of observable harms. In addition, the harms of targeted advertising could occur in many contexts, especially those that promote harmful social biases. In summary, it seems that *non-maleficence* is primarily context specific to social media and health app harms; broadly associated with targeted advertising harms; and what harm consists of is mostly universally understood.

(2) *Security*

Security was closely related to *non-maleficence* but focused on *protecting oneself from harm* rather than others. *Security* was primarily related to protecting one's sensitive data from harmful, unintended, third-party uses. It was usually discussed in terms of data protection, data leaks, and (government or law enforcement) surveillance. *Security* was a prevalent sub-value throughout the data and often motivated specific privacy decisions. Nearly every participant expressed some kind of sensitive data that they never shared due to the risk of it being used by unintended parties. Sensitive data included relationship status, financial information, phone number, intimate photos, or location. Data collected by voice assistants such as Alexa was also considered to be of a possibly sensitive nature that could get into the hands of unintended third parties.

P09: “[...] stuff like Google Home and Alexa [...] being subpoenaed by [...] different judicial and law enforcement authorities. That's terrifying. [...] We welcome the potential for interference into our homes [...] and then we act surprised when we get interfered with.” [Emphasis added]

As P09 demonstrates, participants expressed fear of their sensitive data ending up with law enforcement. Other parties of concern included (particularly the US) government¹⁵⁹ and criminals who may cause them harm. Participants were especially concerned about their sensitive data being used for identity theft, extortion, scamming, and stealing money from their bank account. P15 even shared their experience of having their passwords stolen and being used to extort them, and the effects it had on their feeling of *security* (Q15, Appendix IX). For another participant, *security* meant not being contacted by unknown persons using

¹⁵⁹ Likely a reference to "The Patriot Act" or similar surveillance activities in the US (see Section 2.2 or footnote 163), popularized in the EU following the Schrems I and II rulings by the European Court of Justice. In summary, these rulings established that data transfers to the US from the EU were in violation of the GDPR and other EU regulations because the data could be accessed by US intelligence agencies. To redress these concerns, an agreement between the EU and the US was reached, and President Biden signed an executive order to implement additional safeguards in 2022. For more, see: https://commission.europa.eu/law/law-topic/data-protection/international-dimension-data-protection/eu-us-data-transfers_en

How Do We Value Data Privacy?

their shared data, such as their phone number, in an attempt to scam them or get them to purchase a product (Q16, Appendix IX). Others noted how the online sphere is ripe with opportunities to be scammed. Some of these scamming attempts are more benign, perhaps even comical, and easy to identify, while others utilize “finely tuned dopaminergic stimuli” to draw us to give away our information (e.g., “romantic scams” (P08)).

While *security* was primarily related to protecting one’s sensitive data from unintended and particularly harmful uses, it could motivate alternative actions depending on the individual. On social media, one participant avoided tagging people in photos, only sharing photos years after they were taken, and making sure to limit their social media network to close family and friends. Other strategies included avoiding services with long terms and conditions, viewing this as an indicator of greater data protection risk, and only uploading sensitive documents on their computer and not their phone, which was viewed as less secure. Pivoting to app choice, two participants avoided bank apps or Google Pay/Wallet due to concern over stolen financial information. Backdoors for unauthorized access to smartphone apps were also a concern. P26 made an analog between the security of the apps you keep on your phone and your home, viewing *security* as central to privacy.

P26: “I think privacy is really important to people. I wouldn't like somebody just to walk into my house and start snooping around. Same thing with my bank account. You know, same thing with my phone, you understand? [...] It's those kinds of things that that create problems. [...] if you leave your front door open, somebody can just walk in, right? If you if you allow a backdoor to be opened [on] your phone, you're pretty much saying to people: “help yourself! [...] Yeah, come on in! Take a look around!” [...] So privacy is really, really important.” [Emphasis added]

There is some similarity here with the “digital home” analogy in Chapter 2, where our apps and the data we share with them are like furniture or people we allow into our home.¹⁶⁰ In this case, P26 aims to secure their “digital home” from unintended visitors, because they value *security* and do not want someone coming in and “helping themselves” to their personal data. Relatedly, another participant described following “technical letters that come to [their] email every month” that identify apps with security vulnerabilities or questionable data collection practices (P15). Besides avoiding or deleting apps, *security* also motivated downloading new apps. Two participants felt that an app is secure if it has many downloads (e.g., users), viewing this as an indicator that the app has a robust data protection policy. This motivated them to choose mainstream apps over other less popular apps. Thus, not sharing sensitive data to promote *security* also translated to not engaging with apps that require such data or appear to have a risk of unintended data collection or use. This result – a value like *security* motivating specific app choices – supports our understanding of app choice as a critical privacy decision point, where we have an opportunity to exercise some control over what data we share by selecting a value-consistent app (Section 2.5).

While sharing data was frequently associated with less *security*, this was not always the case. Two participants felt that there are instances where sharing sensitive data like location can actually promote *security*. P12 discussed how location sharing can allow people to “actually see registered sexual offenders near [their] area so that [they] can be much more aware of those things.” P30 also mentioned that sharing location data on their phone makes them feel more secure in case something malicious was ever to happen to them, such as getting “kidnapped. Maybe it's good that [their] phone knows [their]

¹⁶⁰ See footnote 140 or Section 2.4 for more on the “digital home” analogy.

How Do We Value Data Privacy?

location.” Sharing data when reasonable protections were in place was also not viewed as undermining *security*. One participant expressed that companies have data protection policies to protect the data they collect, and this makes them feel secure enough to share their data with them. Another expressed how encryption and 2-factor authentication made them feel secure enough to share financial or other sensitive data with an app. In this case, they did not feel that sharing data with such protections in place violated their value of *security*. This suggests that individual motivational relationships between *security* and one’s actions are possible depending on whether the individual feels the data will be used to promote their *security*.¹⁶¹

In summary, *security* was primarily related to protecting one’s sensitive data from unintended and particularly harmful uses, including not engaging with apps. This suggests that *security* is largely related to (restricting) sensitive data-sharing, essentially regulating access to one’s “digital home.”

(3) Trust and Trustworthiness

Participants also expressed the importance of *trust and trustworthiness* when it came to data sharing, which was intimately related to their sense of risk and safety (*security*) when sharing data. *Trust and trustworthiness* was expressed in terms of choosing an app, using indicators that the app was sufficiently secure and worthy of entrusting their data to. The most prevalent indicators of *trustworthiness* were that the data the app requested matched with the app’s function; the number of people using the app, with higher usage an indicator of greater *trustworthiness*;¹⁶² the presence of positive app reviews; and whether contact info for the app was provided.

While there were these shared indicators of trustworthy apps, indicators more commonly varied considerably based on the individual. To begin, three participants expressed that receiving an ad for the app made them trust it less.

P02: “So I think the Calm app comes up a lot on [...], Instagram. That actually like disincentivizes me personally, because [...] when I start getting [it] advertised [...] that's when I start thinking like, "Ah, you're already probing too much into me." [...] That's when I start to distrust things [...] [they are] trying to get money out of me or induce me.” [Emphasis added]

In contrast to P02’s Calm example, one participant expressed that apps that were informed by experts (e.g., Headspace) made them more trustworthy. In addition, an app that has never given participants any reason to distrust (e.g., data leaks or causing them harm) increased trust for some. A few others mentioned other indicators of trust, including: 1.) persistent app notifications, deemed less trustworthy because it suggests the app does not have your best interest at heart; and 2.) whether the app was recommended by a trusted third party, such as a favorite podcast.

¹⁶¹ Interestingly, no participants mentioned government surveillance as a means to promote *national security*, focusing instead on *personal security*.

¹⁶² While finalizing this thesis, Maseeh et al. (2023) published their work investigating user privacy concerns, also using Reflexive Thematic Analysis (Braun & Clarke, 2022). They found that an app’s creditability – its “trustworthiness and believability” – was an indicator of privacy concern, willingness to download an app, and willingness to share data with an app. While their research aims and theoretical lens were different from what is used in this work, this overlap with our sub-theme *trust and trustworthiness* suggests a strong influence of app popularity when it comes to app data privacy decisions.

How Do We Value Data Privacy?

While *trust and trustworthiness* was predominantly discussed in terms of downloading/choosing an app, it was also discussed with similarly varying indicators when *trusting* online services. In these cases, long terms and conditions were viewed by one participant with suspicion (P26: “They have to be fancy about it. And they can’t be upfront [...] That’s really disconcerting.”) For another participant, in contrast, these long terms and conditions inspire a (possibly false) sense of *trust* (P15: “[...] whether it’s effective or not, it’s giving me a comfort [...] psychologically”). Despite widespread variation, there was one consistent indicator of service *distrust*. Whether the service has a history of data misuse was also a motivator for trust, with Meta/Facebook frequently viewed by participants as untrustworthy for this reason (Q17, Appendix IX).

P01: “I do not think that Google has been found out as Facebook has to be selling the data to political agents as clearly and as target focused as Facebook. [...] I still trust Google a bit more than I trust, well, Meta. [...] I am less trustful of companies like Facebook, for example, that have a record of misusing users’ data for sort of at least gray purposes, if not downright evil purposes.” [Emphasis added]

Relatedly, participants also discussed different entities in terms of their *trustworthiness* based on their own personal set of *trustworthiness* indicators. They discussed companies, governments, and researchers. Besides distrust of big companies due to data misuse, some participants also distrusted companies because they had a profit motive. In contrast, another group of participants expressed greater trust in big companies, because they have the resources to have data protection infrastructure in place and are beholden to laws.

P06: “I think bigger companies [...] usually they’re sitting in a developed country with developed laws, so they’re under certain law, [not] just some developer somewhere in any country. If it’s in my [home] AppStore, I think it’s still under certain laws, but like, if the data goes there, they may be doing whatever [...] but like, a big company, they have some values, some ethics, and [...] hopefully stick to [them].” [Emphasis added]

Trust in governments was similarly mixed. For one participant, governments were viewed as a “lesser of two evils” compared to companies because “at least [governments] theoretically have the citizens’ best interests at heart as opposed to a corporate entity that is just about profit” (P09). However, P09 also expressed some concern that data entrusted to governments could further prop up existing problematic surveillance structures, mentioning surveillance in the US (“The Patriot Act”)¹⁶³ and China.¹⁶⁴ They also referred to “incompetent” handling of citizen data in their country as a possible risk to citizens (Q18, Appendix IX). In contrast to P09, P02 believed that government bodies are more trustworthy than companies when it comes to sensitive data collection, such as health data collection. For them, the lack of profit motive increases government *trustworthiness*, arguing if “it’s a health service app, like, if [governmental health body] is putting out this tracker, they’re incentivization for this app [...] actually to help me with my health.” Lastly, P06 stated that we should *trust* researchers more than companies with our data.

¹⁶³ “The Patriot Act,” briefly touched upon in Section 2.2, was signed by President Bush in 2001 in response to the 9/11 terrorist attacks. It resulted in an expansion of US counter-terrorism surveillance (Ombres, 2015). The extent of this surveillance was famously exposed by whistleblower Edward Snowden in 2013. See *Permanent Record* by Edward Snowden (2013) or *The Guardian* reporting (<https://www.theguardian.com/world/2013/jun/09/edward-snowden-nsa-whistleblower-surveillance>) for more information.

¹⁶⁴ For example, technology has been used extensively to surveil China’s Uyghur minority. See: Leibold (2020).

How Do We Value Data Privacy?

Like distrust in companies, researchers were deemed more trustworthy because they did not have an individual profit motive. Instead, P06 viewed researchers as motivated by a desire to make the world a better place, carefully undergoing “lots of ethics training [and] data training” to “get a sense of ethics” in the pursuit of responsible research.

Taken together, what indicates a trustworthy entity for sharing data – and therefore allows one to act in accordance with their value of *trust and trustworthiness* – is highly personal, although certain common considerations emerged in the research. Which indicators one decides to use also shapes which entities they choose to share their data with.

g Use

The value *Use* was oriented towards the function, usability, or technical specifications of an app or service. It was one of the most prevalent values in the data that motivated privacy decisions – affecting both decisions to download an app and to share data. This suggests that *Use* is particularly relevant to privacy decision-making. There were two sub-values for *Use*: *time and convenience* and *utility and function*. *Utility and function* was one of the most prevalent sub-values in the data and spanned contexts. Indeed, as we will see, the value of *utility and function* is intimately linked to the app or data-sharing context.

(1) *Time and Convenience*

Time and convenience was consistently understood as promoting ease and speed in all aspects of life. Most of the discussion focused on the convenience and time-saving value of smartphone apps. Half of participants mentioned the value of apps for making everyday life more convenient and efficient. Apps allowed participants to save time by writing a thought down in Google Keep (P28); booking an appointment for a “nail bar” using an app (P26); internet browsing or accessing information (P13, P23); health tracking through smart watch pairing that “doesn't require a lot of effort on my part” (P22); “double jobbin” by listening to audiobooks or a podcast while exercising (P08); or communicating and connecting with others on WhatsApp (P01, P06, P08, P23), Instagram (P07), or other social media and messaging services (P01). Apps could also help participants with their time management skills, by helping them focus (e.g., Pomodoro timers), plan, or improve their productivity using “Brain Training apps” (P09). P23 further linked this time-saving convenience of apps to smartphones themselves, describing the convenience of having these accessible, powerful computers in our pockets (Q19, Appendix IX). The role of having data centralized in saving time was also discussed by three participants. Although participants expressed reservations towards data centralization (discussed in *power and choice* and its tensions), they did appreciate the convenience of having their data all stored in the same place. In particular, two participants mentioned the benefits of being able to transfer data from their Google account when getting a new phone.

Participants also mentioned the role of *time and convenience* on their data sharing behaviors, usually associating managing one's data privacy as inherently in conflict with *time and convenience*. This is interesting, as we previously theorized that those who value *efficiency* may become overburdened with PPA notices, introducing a *double bind* (Section 3.3). In this vein, three participants described choosing to take “the path of least resistance” (P14).

How Do We Value Data Privacy?

P14: “Now there might be a negative cost down the line that you're not necessarily seeing [...] but at the time, **when you're signing up to the app** [...] [you take] **the path of least resistance** [...] just go, “Yeah, I agree.” You know, how many of us actually read the terms and services? We just kind of go, “**Yeah, I agree to this, like, blah, blah, blah, blah.** Yeah, I agree. [Laughter].” [Emphasis added]

Another also felt “time and it's a headache” was holding them back from “installing a pure Android operating system to [...] get rid of Huawei's bullcrap” (P01), referring to the default, privacy-invasive, apps that came with their Huawei phone. In a few sections, we will explore this tension between *Use* and *Control* in more detail, as well as consider whether it is a *double bind*.

It is also notable that, while understandings of *time and convenience* were similar, some participants did not value it as much when making data privacy decisions. One participant viewed this taking of “the path of least resistance” when it comes to privacy decisions as an excuse, bringing out our easily manipulated, “comfort loving animal” nature (P08) (Q4, Appendix IX). In summary, *time and convenience* was simply less relevant to data privacy decision-making for them.

Besides these participants, the results overall suggests that *time and convenience* is a broadly relevant value for most users. In these cases, it is consistently valued in privacy decisions to motivate quick, efficient choices.

(2) *Utility and Function*

Utility and function focused on the usefulness, necessity, purpose, and technical specifications of an app or service. This sub-value was one of the most prevalent in the interviews, mentioned by nearly every participant and in many different contexts and situations. Participants discussed *utility and function* within three broad areas: apps as valuable for the *utility and function* they possess; the value of sharing data to increase *functionality* and usability; and disengaging with online services that no longer served their *function*.

Firstly, apps (and some online services) were consistently viewed as modern-day utilities whose necessity was intimately linked to the kind of app. In some cases, the necessity of the app was determined by external forces, such as requiring it for a job. In other cases, the degree to which context-related (sub)values were being promoted were directly related to the value of *utility*. For example, WhatsApp and social media were frequently viewed as necessary utilities for news gathering (*staying informed*) and communicating (*connection*). The usefulness of translation services was also mentioned (*connection*). Participants also valued apps if the function allowed them to pursue their goals, studies, or hobbies, such as hiking (P13), music (P01, P22, P28), or photography (P30) (*authenticity; power and choice; self-improvement*). On the flip side, participants also described the crucial role of social media (e.g., Twitter, Instagram, and LinkedIn) as an essential networking or promotional tool for existing as a modern-day professional (whether they like it or not). Email apps and apps for online cloud services, such as Google Drive, were also viewed as necessary utilities for doing their jobs. When commuting or traveling, apps were again viewed as necessary utilities for accessing local transit services. Google Maps was similarly viewed as a modern-day utility that is essential for getting around in the world, as were apps to manage one's finances.

Besides the necessity of the app, the *utility and function* of the app was also considered in terms of its technical capabilities and requirements. These included WIFI requirements, size and available phone memory, battery life, and interoperability with

How Do We Value Data Privacy?

other devices, such as Alexa. Based on their own situation, different participants considered these technical constraints differently. Considerations about the size of the app ranged from “500” apps on their phone with sufficient memory to spare (P07) to fears of having their phones’ “memory [...] almost full” (P23), the latter having to choose and delete apps with more care. In addition, having apps on their phones that were clunky” or poorly designed was frequently stated as a motivation to delete an app (P07). Depending on an app’s *utility and functionality* alignment with the participant’s interests, three participants further stated that they would be willing to pay for increased *functionality*. For example, P26 described one of their favorite apps, Carb Manager, to help them meet their weight loss goals as being well worth the price for the features and services it provides.

Participants also expressed this value of *utility and function* when searching for these apps in app stores¹⁶⁵ by going in with a specific function in mind and using keyword searches based on those app’s desired function. To assess the value *utility and function* when looking for an app, participants also felt that better reviews or more downloads were positive indicators of *function*.

P04: “I rarely just scroll through the app store. I usually have something in mind that I want to do, or [...] a functionality that I might use. I just recently download[ed] a QR scanner. So, I just type “QR scanner,” look through it, and usually will download the one with the best reviews. [...] So, it's mostly about what I need or what I want to use. But I know that before.” [Emphasis added]

Participants also described initially downloading apps and testing them in terms of *functionality* and usability. Two participants described how they download apps to test and then delete them if they are not functional or useful. For some, *utility and function* was the main value they considered when choosing apps (P06: “My phone is very functional”). Apps that came with the phone that could not be removed but did not serve a function were described as unnecessary “fluff” (P26).

While participants primarily discussed *utility and function* in terms of app choice, engagement, and deletion, the value was also discussed outside apps. Firstly, adding to concerns around data centralization (*power and choice*) and its value to *time and convenience*, participants also noted that “the utility of having [their data] all together is so huge” (P02) (Q20, Appendix IX). Secondly, lacking *functionality* or *utility* also motivated users to disengage with online services. Facebook’s decreasing *functionality* (P14: “90% of the things on your profile were competitions”), for example, motivated some participants to delete their Facebook accounts (Q21, Appendix IX).

In summary, *utility and function* is very relevant to app choices and (more traditional) data privacy decisions. It spans contexts and individuals; it is similarly understood; and it tends to motivate similar decisions – everyone, after all, wants apps and services that are *functional* and useful, with some even prioritizing this sub-value above all others. However, *utility and function* is intimately associated with an app or service – for example, *connection* and WhatsApp. This is an interesting finding, as it suggests that how well *utility and function* is promoted by an app choice or privacy decision is intimately linked to how well other context-dependent (sub)values are being promoted.

¹⁶⁵ How participants searched for apps in the Mock App Store and the implications of these search strategies for VcPA design are discussed in Chapter 6.

h Value Conclusions

The value themes described here suggest a complex, motivational relationship between a user's values and privacy decision-making. The value of *Control* was quite prevalent in the data. Besides *Control* (especially *power and choice*), *Use* (*utility and function*), *Community* (*connection*), and *Safety* (*security*), were also quite prevalent. Many of the values tended to gravitate towards specific contexts or situations – for example, *Community* towards social media and *Safety* towards sensitive information. In addition, different understandings of values and how they could best be acted upon led to largely varied, individualized,¹⁶⁶ relationships between values and privacy preferences. The only exceptions to this were *Control* (especially *power and choice*) and *Use* (especially *utility and function*), which, for the most part, spanned contexts and were similarly understood. Interestingly, how well *utility and function* is promoted is intimately linked to how well other context-dependent (sub)values are promoted, such as *connection* and WhatsApp.

These results also have implications for our understanding of values in privacy decision-making and for the 4DT conceptualization of value-centered privacy decisions presented in Chapter 3.

Firstly, emphasis on *Control* (or lack of it) and orientations towards commitments and motivations, for both more traditional privacy decisions and app choice, are consistent with the 4DT-based understanding.

Secondly, *dignity and respect*, also under *Control*, shows that users also see themselves and others as individual agents deserving of respect. There is also a sense of responsibility to oneself and one's actions, whether that be to meet goals in order to improve oneself (*self-improvement*) or fight a sometimes-fruitless battle to manage one's data privacy. *Dignity and respect* and this sense of personal responsibility resonate with the self-governance basis of 4DT.

Thirdly, some expected phenomena based on the 4DT understanding of value-centered privacy decision-making (see Figure 2-1 and Table 3-1) were observed, albeit with some added insights that were not captured by our initial conceptualization in Chapter 3. Again, under *Control*, we saw a loss of agency under existing power structures, resulting in an inability to act according to personal *power and choice* when making data privacy decisions. In some instances, these external power structures limited choices available to users (lack of controls in a 4DT understanding), meaning that participants were unable to *self-realize* and *self-unify*. In addition, we observed the “apathetic user” phenomenon, where a user fails to *self-constitute* because of becoming apathetic in the face of such power structures. While the “apathetic user” was initially understood in terms of third-degree notice fatigue, the influence of power structures is harmonizable with a 4DT approach to value-centered data privacy decision-making. The design choices described by participants to force their data privacy decisions were also in accordance with the 4DT (and most other) understanding(s) of dark patterns as autonomy-violating (Chapters 2 and 3). The *inertia bias* was also present, but interestingly, applied to keeping “old” apps as much as sticking with the defaults of data privacy decisions, although, again, this would not be inherently inconsistent with 4DT and can be understood within his theory. There were also instances of *weakness of will* when making data privacy decisions. *Time and convenience* under *Use* showed us that some users indeed find too many privacy notices problematic, as expected when we evaluated PPA notices in Chapter 3.

All that said, the interview data also suggests that there may be a simpler means of capturing and communicating the role of values in data-privacy decisions. Recall from

¹⁶⁶ E.g., the *Community* sub-value of *accessibility* only expressed by P26, or sharing location data viewed as protecting, rather than undermining, *security* (under *Safety*) for P30.

How Do We Value Data Privacy?

Chapters 2 and 3 that 4DT was selected in part for its granular dimensions, which allowed us to design a value-centered privacy assistant. While it is promising that the theory we utilized to identify these design features appears to be able to capture relevant phenomena, the added complexity of a dimension-by-dimension analysis may not be necessary when communicating to other stakeholders, computer scientists and policymakers. For example, consider that we wish to communicate how users are struggling to make value-centered privacy choices according to their valuing of *power and choice*. We could discuss the problem in terms of 4DT and its dimensions, pointing out that it was specifically the dimensions of *self-realization* and *self-unification* that were being violated. However, little would be lost if we instead described how users value *power and choice* over their data privacy decisions and that existing power structures limit their ability to make the choices they wish, thereby violating their autonomy.¹⁶⁷ In sum, while the dimensions of the 4DT were critical from a design perspective for identifying necessary features to promote value-centered choices, the level of granularity may not be necessary to capture and communicate the role of values in privacy decision-making.

Taken together, the interview results suggest that the 4DT understanding of values in privacy decisions can reasonably capture the role of values in privacy-decision making, although all its dimensions may not be necessary when communicating the role of values in privacy-decision making to relevant stakeholders. In most cases, how *specifically* a particular value is related to privacy decisions is complexly driven by both an individual's understanding of the value and the context in question. However, two values – *Control* and *Use* – were similarly understood and relevant regardless of individual and context.

5.3.4 Value Tensions and Resolution Strategies

a Value Tensions: Overview

In total, fourteen tensions between the six values were coded in the interview data (Table 5-3). There were five tensions that appeared in more than five interviews, all involving either *Control*, *Use*, and/or *Community*. The other nine appeared less frequently, sometimes specific to one participant. While most tensions were between values (“inter-value” tensions), three of these infrequent tensions were within values and between sub-values (“intra-value tensions”). Participants also frequently described ways of resolving tensions. However, resolving the major tensions (between *Control*, *Use*, and *Community*) was generally perceived as difficult, if not impossible.

These results suggest that resolving tensions between *Control*, *Use*, and *Community* could be the most critical for promoting value-centered privacy-decisions. This is further supported by the prevalence of these values which, as described in the previous section, were frequently expressed by participants. We also identified a spectrum between resolvable tensions and “true” *double binds* depending on the participant's values and their ability to act. This suggests opportunities to support users' existing resolution strategies in the pursuit of more value-centered privacy choices. However, the concept of *double binds* was helpful to identify the most critical areas where users may not be able to act in full accordance with their values – that is, instances where the tensions are most “*double-bind* like.” In these cases, the presence of structural factors – such as social media monopolies, the attention economy, and surveillance capitalism – were defining features. In the coming

¹⁶⁷ Here is another example, also from *power and choice*: we could describe the observed “apathetic user” phenomena brought on by existing structures in terms of a decrease in self-governance, or autonomy, rather than *self-constitution*.

How Do We Value Data Privacy?

pages, these insights will be discussed in more detail within the context of each of the fourteen tensions.

Table 5-3: Summary of value tensions, with three tensions of particular interest in bold

Prevalent Tensions	Minor Tensions
Control-Community	Inter-Value
Control-Use	Control-Growth
Community-Use	Growth-Community
Safety-Community	Growth-Use
Pleasure-Use	Pleasure-Control
	Safety-Growth
	Safety-Use
	Intra-Value
	Community (authenticity-connectivity; conformity-benevolence and universalism)
	Use (time and convenience vs. utility and function)

b *Prevalent Tensions*

(1) *Prevalent Tensions: Overview*

The most frequently discussed tensions were *Control vs. Community*; *Control vs. Use*; *Community vs. Use*; *Pleasure vs. Use*; and *Safety vs. Community*. While tensions were largely individualized and context-dependent as seen with the values, the most prevalent tensions involving largely the same values suggest that these values could be the most critical to resolve to allow for value-centered choice. Participants also presented their strategies for resolving these tensions (upholding *self-realization* and *self-unification*). They also expressed how they could be further supported in resolving these tensions, with tensions existing on a spectrum between generally resolvable tensions to “true” *double binds* depending on the participant’s values and their ability to act. This suggests opportunities to support exiting resolution strategies to promote more value-centered privacy decisions. However, resolving the major tensions (especially those between *Control*, *Use*, and *Community*) was generally perceived as difficult, if not impossible, suggesting these tensions were more like *double binds*. This was largely due to not having an avenue of action that would be consistent with both values, usually as the results of structural barriers.

(2) *Control vs. Community*

Control vs. Community could be best summarized as a “love/hate” (P22) relationship with *Community*-promoting apps and services and the limited ability to say “no” to them (*Control*). This was especially prevalent with social media or messaging services, such as WhatsApp. WhatsApp was seen as valuable to *connect* (sub-value of *Community*) with others, but participants also felt they had little choice in having it (*power and choice*).

P08: “I feel like everybody's online now, in some form or another, and because I don't do social media [...] I feel like I need to be on WhatsApp [...] not many people meet up and chat like they used to.” [Emphasis added]

How Do We Value Data Privacy?

Sub-value *conformity* (under *Community*), in the form of peer pressure, also came into tension with *power and choice*. In these cases, going along with what others are doing was at odds with one's own concerns about their engagement with technology. Limiting social media based on one's choices was viewed by one participant as antisocial, and four expressed embarrassment or doubt as to the legitimacy of their data privacy worries (P26: "Look I'm not completely cuckoo. I'm slightly cuckoo [joint laughter]"). Six participants felt perceived by others or perceived themselves as "old school" if they did not accept this Brave New World(s) of data collection and apps (P13). While they would currently "never have [app] downloads, for example, [like] Google Pay, [...] you should do this interview in three years [...] and see where [they are at] then" (P13). They felt that they might have to change at some point to keep up with the times. This pressure to conform could also work in reverse, with peers encouraging more data privacy-preserving actions. One participant expressed embarrassment for their more relaxed views towards data privacy because their friends frequently discuss the data protection risks in their jobs.

Uniquely, one participant also felt that people are usually too worried about researchers' collecting and using their data, where the good of the research (*benevolence and universalism*, under *Community*) can be interpreted as conflicting with *tolerating others' choices* (*tolerance*, under *Control*). There was also tension between *benevolence and universalism* and another sub-value of Control, *power and choice*, when it came to how data was used in targeted advertising. One participant noted the benefits of targeted ads for finding products they like (P09, *benevolence and universalism*), while also noting that there is not currently enough user control over them (*power and choice*). Another noted that targeted ads can help them find apps they would like to download but also could use their wellness goals to "trap" them into downloading an app they do not need (Q27, Appendix IX).

To manage this tension between *Community* and *Control*, participants described several strategies that varied based on the individual and the sub-values involved. To resolve the tension between *connection* and *power and choice*, one participant decided to only use social media on their computer because the app kept asking them to share data (their phone number) that they were not comfortable sharing. This participant also keeps their social media anonymous to protect their privacy, using a pseudonym. Another participant decided that they don't really care if some types of data or information are available on social media (P30: "[...] if somebody wants to look at pictures of my cat, more power to them!"), but is mindful of sharing more personal information or photos, such as selfies. Indeed, limiting one's social media presence in some shape or form (e.g., limiting social media accounts, posts, photos, visibility with a private account, or only "friending" close friends or family) was the most frequent strategy to manage this tension.

Managing this tension using these strategies, however, was not perceived by all participants as satisfactory for resolving the tension between *Community* and *Control*. Participants also noted that regulation to stop ad retargeting (P09; resolving *benevolence and universalism* vs. *power and choice*) and a decentralized network for social media to keep their data sharing to close friends and family (P10; *connection* vs. *power and control*) would support them in resolving this tension. Two participants also noted that attempts to resolve the tension between data privacy concerns (*power and choice*) and connecting with friends (*connection*) by moving to another platform was difficult, if not impossible, because their networks would not move with them.

How Do We Value Data Privacy?

P01: “I understand that **I could become an activist and ask everyone to contact me through Signal** and that way, I could spread something good. But I don't feel like I have that sort of force in me to do that to, to all of my friends and so I don't do it. So, I have to [have WhatsApp]. **The only solution is to comply.**” [Emphasis added]

Situations such as these are indicative of *double bind* situations where the agents, such as P01, cannot act in full accordance with their values – but must pick one route to follow. This is, therefore, an ideal situation for promoting more value-centered choice, designing in a manner that allows users to *self-realize* and *self-unify* in accordance with both, *power and choice* and *connection*. As discussed when evaluating the PPA using 4DT, one means of overcoming this would be to suggest alternative apps. However, like the case of Shauna (Table 3-1), who wishes to choose a similar app with similar *connection* value to WhatsApp, such an alternative may not exist. This is because we are constrained by where those we wish to connect with are and the market dominance of the app in question. Similar to promoting *power and choice* alone, resolving the tension between *power and choice* and *connection* will likely require regulatory or more global interventions that promote one's ability to *self-realize* and *self-unify*. In particular, social media monopolies could be broken up to increase the availability of viable alternatives.¹⁶⁸

(3) Control vs. Use

Control vs. Use usually captured the feeling participants had that they would rather not share their data or have this app (*power and choice*), but that it was necessary to do so because of the utility provided (*utility and function*). This was usually expressed by guilt and shame around sharing data or engaging with social media when they know it is not in their best interest to do so.

P07: “I don't let [my data privacy concerns] you know, stop me from downloading anything because **I don't care enough**. Whereas I do know people, you know, like that would say, “Oh, well, this apps tracking too much stuff. So, I'm just I just don't need it” [...] I guess maybe I don't care enough. Like, **I feel like I do care**. [...] I feel like I do have a strong opinion. Like, **I don't think apps should track you**. I don't think they need that information. But it's not I guess it's [not] enough for me to say, “Well, I'm not going to use the app” because, like, **I still want to live my life.**” [Emphasis added]

This sentiment to “live one's life” captures the necessity of these apps or information sharing to exist in the modern world while still wishing for more control. Apps have become necessary¹⁶⁹ to access critical, everyday services, such as transit (e.g., planes, taxi, bikes) or bank apps (Q22, Appendix IX). Limited functionality or availability of alternative services also pressured participants into downloading apps. One participant tried managing their bank account via the website, but “found it really difficult to use,” adding: “they really pushed you to use the app” (P13). Participants also mentioned “there's usually not too much flexibility in what [data] you can limit” if you wish to download an app or access a website (P07). When these apps or services are needed, then participants expressed that they really did not feel they had much of a choice because “it's difficult to [...] keep following your values when you have to download apps that are necessary for [wherever]

¹⁶⁸ To be explored in greater detail in Chapter 7.

¹⁶⁹ While not always viewed as a necessity per se, the utility offered by storing data in a central service, for example, Google services, was also expressed to be in tension with one's control over their data (Q20, Appendix IX).

How Do We Value Data Privacy?

you're traveling.” (P13). There is little choice but to trust these apps with their data and wish for the best (P07: “Well, hopefully they don't misuse the data. But here you go.”).

While most of the *Control* vs. *Use* tension was exhibited between the sub-values *power and choice* vs. *utility and functionality*, there was also tension between *power and choice* and *time and convenience*. Some participants expressed that managing their data privacy, whether it be on their smartphone or online, was not always feasible and would take too much *time*. When prompted with a privacy notice, this *time* concern would encourage them to quickly click “Accept All” or go with the default settings (Q23, Appendix IX).¹⁷⁰ This was also true when it came to the default apps that come with, say, a Huawei phone.

Participants also described strategies for managing or resolving this tension. Nearly every participant resolved the *Control* vs. *Use* tension with a strategy of data privacy pragmatism. Reminiscent of Westin’s data privacy pragmatists (Sections 2.2 and 3.3), data privacy pragmatism here is understood as a cost/benefit analysis, a ranking of *utility* of the app or the *time* required to manage data privacy against one’s *control* over their data, individualized to each participant’s particular needs and data privacy concerns. Different aspects associated with *Use* and *Control* were given more weight depending on the participant, with most focused on the *utility and function* value of an app or service against *power and choice*. Most used an approach of “function first, data privacy second” to manage this tension, willing to go with an alternative app that better matches their privacy preferences if it also had the desired *functionality*.

*P02: “I want to be able to compare the utility against the cost [...] utility of the functionality of the cost of privacy. And then I get there was **very little difference in the functionality**, then yeah, I probably pick the one which gave me a better privacy. But if it wasn't giving me the functionality, I think **I'll just drop the privacy pretty quickly.**” [Emphasis added]*

Many participants also added that an important consideration was whether the data collected is necessary for the app or service to function. If so, participants considered such collection acceptable, and were willing to forgo some *Control* over this data for the *utility* or service the app provided. Otherwise, more weight was given to *controlling* data that did not seem to be necessary for the app or service to function, motivating participants to turn off certain permissions on their smartphones or not downloading the app altogether. Three participants further embraced this pragmatism by installing and deleting apps as needed as a means of balancing *Control* over their data and app downloads against the *utility* of the app or service.

*P14: “So, one [night I was] coming home. [...] I couldn't flag down a taxi. So, I installed Free Now and ordered through that, got the cab home, deleted it straight away. You know, Ryanair would be similar. **If I'm not actually going flying anytime soon, I delete it off my phone as well.**” [Emphasis added]*

Other examples of pragmatic strategies included: accepting (most) data sharing in favor of *utility* if the data shared is associated with a user ID and not their name (“anonymous”) (P06) and limiting public features in order to balance the *utility* of the app or service against *control* over their data. For example, P14 stopped using Goodreads “to friend

¹⁷⁰ It is notable that one participant felt that statements such as these were just an excuse, and we just need to “pay attention” and do the “two second job” to manage our privacy (P14) (Q24, Appendix IX)

How Do We Value Data Privacy?

anyone or import contacts” to “use it solely as a book tracker, because [they] want to give away as little information as possible.”

However, misunderstandings about privacy, what data is being collected, and how it is being used may mean that such data pragmatism strategies are not truly resolving the tension (*self-unifying*). Misunderstandings, some of which have already been presented in terms of values and value tensions associated with them, were: closely associating privacy with data protection and scams which, while related, are not the same (7 participants); believing the (Apple) Appstore (always) vets apps to make sure their data collection practices are safe, secure, and transparent (1 participant) (Ali et al., 2023; Jain et al., 2023; Koch et al., 2022; Kollnig et al., 2022; Paci et al., 2023; Rodriguez et al., 2023); the “I’ve got nothing to hide” argument when it comes to data sharing (Solove, 2007); “my phone is listening to me” when getting a targeted ad after talking about something (Khan, 2021); that paying for an app (always) means less data collected (Brumen et al., 2023) (2 participants); and that you need to accept all cookies to access a website (one participant). These misunderstandings, in addition to the previously-well documented lack of user awareness of data shared and understanding of controls,¹⁷¹ could also hinder users from making privacy decisions in accordance with their values by depriving them of the accurate information needed to resolve tensions, such as the “function first, privacy second” strategy to resolve *Use vs. Control*. In reality, they are not resolving this tension at all – they are likely acting against *Control* based on false information. Here, though, we risk falling into the impossible challenge of fully informing the user – the “transparency paradox.”¹⁷² While the value-centered approach (ideally) helps with this literacy challenge by shifting the focus to the values behind data-privacy decision-making, the link between values and privacy preferences is still something that needs to be clear to participants to be effective.¹⁷³

While the data privacy pragmatism was the most prevalent and the most interesting when considering our 4DT understanding of values and data privacy decisions, a smaller number of participants utilized certain self-imposed rules, rather than pragmatic strategies, to manage the *Control vs. Use* tension. For one participant, they immediately deleted apps that require any kind of subscription or payment (P29). Others added more weight to sharing certain kinds of sensitive data which, if asked by the app, trumped all *utility* of the app or service.

*P09: “If [the app requires] sharing my email address or home address, [...] then **it doesn't matter how good the app is**. That's very sensitive information.” [Emphasis added]*

To maintain some *control* over their data while still accessing a website, P30 decided to instigate a policy of accepting all cookies by “install[ing] “I Don't Care About Cookies”” on Chrome. However, they also “do clear [cookies] quite frequently.”

Regardless of the approach, participants suggested improvements to further enable them to apply their strategy of choice to manage the *Control vs. Use* tension. Privacy notices for an app or website should be clearer regarding what data is required for it to function and what is not. The default should be only necessary cookies or automatically “select their preferences, which are already to that standard” (P13). To promote greater control via smartphone app permissions, P22 also thought apps and services should make it easier to not accept certain permissions instead of being forced to take an “all or nothing” approach.

¹⁷¹ E.g., see Turow (2003) . Also nicely summarized in Solove (2021).

¹⁷² See footnote 20.

¹⁷³ This is discussed in greater detail in Chapter 6 when discussing the VcPA.

How Do We Value Data Privacy?

P22: “[Websites could improve by] making it easy to not accept things, because I feel like websites make it so hard to do that. Or [an] app. [...] [and also], if you don't accept it, you can [still] use the app.”

This could also be a *double bind* situation, resulting from a lack of relevant controls – although, in this instance, P22 did not seem too bothered by having to choose one or the other. Rather, they felt that more choices would be better than not being able to engage with the service at all. This suggests that, on the spectrum from resolvable value tension to a *double bind*, P22’s choice to engage with a website without their preferred permissions was more of a resolvable tension than a *double bind*.

In summary, the tensions between *Control* vs. *Use* represent a lack of relevant controls and, in some cases, have features of *double binds* – where one cannot fulfill *self-realization* and *self-unification* because they cannot act according to their values (Chapter 3, Table 3-1). Some can realize their values and resolve the tensions, while others lean more towards a *double bind* situation – no matter what choice they make, they will be acting against their value of *Control*. Those who could resolve this tension were able to find appropriate avenues of action, with many using a strategy of data privacy pragmatism. In these cases, their tensions were not very “*double bind-like*,” resembling more resolvable tensions than *double binds*. However, acting according to what they pragmatically intend to do (their *external self-realization*) could be better supported.

(4) Community vs. Use

The *Community* vs. *Use* tension arose primarily between sub-values *connection* and *time and convenience* where social media apps were described by participants as taking too much *time* from other aspects of their life.

P02: “I just think [messaging apps can be] like **a continuous distraction**. [...] I like the idea of trying to be more focused and just doing what I'm doing [...] **[when] I'm out with a friend, I want to spend my time with a friend**, or at [another] moment, I want to work. I just want to better focus on this, trying to have a more like asynchronous life rather than this thing continuously like bombarding me. [...] [When this happens], I don't think [my phone is] enriching for my general life.” [Emphasis added]

Conversely, participants also mentioned how personalized social media feeds can cause different realities that make *connecting* with others difficult. While *convenient*, they don’t facilitate bonding with those in our community around a shared newspaper (Q25, Appendix IX). One participant also expressed the *Community* vs. *Use* tension in terms of sub-values *authenticity* vs. *time and convenience*, where social media and messaging services were viewed as *convenient* but less *authentic* modes of connecting with others. They describe the value of connecting with others more *authentically* using letters instead of a quick text message.

P12: “**I do still write letters, just get in touch with my old being**, but it's more kind of a relic. So, for me writing letter is kind of emotional, but at the same time, **it's the relic importance**. But **the ease, the effectiveness, of communication** is always higher within the technology factor. [...] **if I could replace any of the technology with “old school style,” definitely, 100%, I [would] do that in a blink of eye**. [...] But the thing is that, when you're writing a letter, there is [...] emotion attached to it. Like, I don't know, if you have read a

How Do We Value Data Privacy?

*letter in the past few years or so [joint laughter, inaudible] [...] there's something with that handwriting, and you know that it's not copied from [anywhere online] [...]. It's not a forwarded message. **So, the person who actually wrote to you sat down, took their time, I value time more than anything else, [took] their time, thought of all the things, and it's not editable as well.** [...] when you write it in wet ink, you either have to strike it off, or you have to write a new letter. So, every word is a thought [...] behind those words. [...] [it] allow[s] emotions [...] even if that's a scolding letter, or a hate letter. Still, it's very important because it's very personalized. **So, it's getting all the advantages of technology without having any of the technology.** [...] It speaks volumes [more] than a voice note. [...] Obviously, you can get a voice note and voice notes can still feel very plain. But when you're reading a letter, it's just like reading a book that whatever your thoughts are, you can actually reframe that and read it in your own tone, which actually pleases you. So that's something unique. And you know that it's kind of very unique. **It's an NFT¹⁷⁴ in the real world [...]**" [Emphasis added]*

To manage the tension between *Community* and *Use*, participants described several distinct strategies. Three participants describe deleting social media apps that were distracting or taking too much time from other aspects of their lives. To minimize distractions, P02 and P06 utilized a policy of app minimalism, trying to only keep social apps on their phone that they felt were absolutely necessary for staying in contact with others. P02 also balanced *connection* with *time* concerns by continuing to access some services, where possible, on the computer rather than on their phone.

P02: "I [...] only use Twitter like on my laptop once a day, like in the morning, and check Twitter to see what's going on in the academic world. [Same with] LinkedIn."

Other strategies included putting distracting apps on a separate device and limiting their presence and engagement with social media. For example, P28 does their best to be selective on who they follow on Instagram, and wished Instagram would stop recommending new people for them to follow.

These strategies, however, focused primarily on resolving the tension between *connection* and *time and convenience*. P12, who described the *authentic* value of letters, was able to resolve this tension by choosing moments to write letters and other times to engage in messaging or social media apps. However, the conflicts between *authenticity* vs. *time and convenience* that results from personalized social media feeds causes different realities that make *connecting* with others difficult and did not seem to be resolvable. This has features of a *double bind* situation – that is, is quite “*double bind-like*.” This suggests that moving away from personalized newsfeeds, the cornerstone of attention economies, may be necessary to promote value-centered choices in this instance. This is discussed in terms of the *Community vs. Safety* tension in a few sections.

(5) *Pleasure vs. Use*

Like *Community vs. Use*, *Pleasure vs. Use* also involved concerns around excessive *time* spent engaging with technology. In this case, however, the value of such interactions is not to *connect* with others, but for *Pleasure*. Six participants expressed this tension in a variety of contexts, with social media (e.g., Tik-Tok), YouTube, game apps (e.g., Sudoku), and

¹⁷⁴ Non-fungible token

How Do We Value Data Privacy?

gaming platforms (e.g., PS2). P28 shared their experience with a particularly fun, but *time-consuming*, app, 2048.

P28: “Occasionally, I might download Sudoku. And then I realized that's all I did in the day, and I delete it again. [Joint Laughter] [...] [there are also] those kind of silly puzzle, puzzle-y type apps. [...] I guess 2048 or, you know, it's those kinds of things. It's those kinds of [...] [that] you know, kind of look up and you didn't realize it was dark outside.” [Emphasis added]

Participants proposed a variety of strategies to help manage this tension, with participants by-and-large finding the means of managing the *Pleasure vs. Use* tension. For some apps, P28 described a strategy of download, delete, repeat: “I'll do it to a point and indulge myself and then take it out in my life again for a while.” A few others described keeping their phones very “functional” (P06), with only “one game [app] that I regularly play.” Four participants described certain instances of resolving the tension by fully choosing *time and convenience* over *Pleasure*, “ruthlessly” removing entertainment apps (e.g., social media, games) from their phones that they felt they were spending too much time on (P28). To this end, P06 describes “delet[ing] Instagram a couple of months ago,” because they were “concerned about the amount of time” spent on it. Two participants also described moving distracting apps to another device, such as an iPad. One hypothetical strategy, however, was more drastic; P08 recommended using generative AI¹⁷⁵ to create personalized images that are “quite visceral,” designed to “work for that specific person” such as an “artistic representation” of you “wasting away into your phone.”

(6) *Community vs. Safety*

Community vs. Safety primarily focused on the tension between destabilizing society (sub-value: *security*) and the value of connecting (sub-value: *connection*) with others on social media. Five participants expressed concerns that “echo chambers” (P01, P07, P30) were destabilizing society. P14 also mentioned that “social media has had such a detrimental effect on the world in things like politics and misinformation.” P30 represented greater participant concerns for *security* by describing the Russian social media misinformation campaign to influence the US 2016 presidential election.¹⁷⁶

P30: “The echo chamber thing is definitely dangerous. And we hear about [...] like on Facebook, there are Russian operatives who were creating events claiming to be made by Americans, but they weren't really made by Americans. [...] [there] was no actual person leading this event, but they knew what kinds of people to target to invite to this event to create discord within society. [...] It's freaky. It's nuts to think about.” [Emphasis added]

There were a few participants for which the tension between *Community* and *Safety* manifested between different sub-values. One participant described a tension between *benevolence and universalism* and *non-maleficence*. This participant, who had professional experience in the advertising space, described helping a friend get their small business off the ground using targeted ads while also expressing concerns that targeted ads can be based on potentially harmful stereotypes (Q26, Appendix IX). For two other participants, the sub-value tensions were between *connection* and *non-maleficence*. These participants were

¹⁷⁵ Generative AI, such as OpenAI's ChatGPT (<https://chat.openai.com/>), are AI systems that generate content, including text, images, and videos.

¹⁷⁶ *Operation Secondary Infektion*, investigated at length here: <https://secondaryinfektion.org>.

How Do We Value Data Privacy?

concerned about the mental health harms of consuming social media, whether it be the constant stream of “bad news” (P07) or getting into online arguments.

P14: “I had a had a Twitter account at one stage that I deleted [...] you're getting involved in interactions probably arguing with people on Twitter a lot and it got to the stage where that was that was no longer enjoyable.”

Still, each felt that there was value to social media for connecting with others, especially for those who lived away from family.

To manage these tensions between *Safety* and *Community*, participants proposed mostly unsatisfactory¹⁷⁷ management strategies and more frequently pointed to ways social media could be improved. This leans towards a *double bind* situation – that is, is quite “*double bind-like*” because the tension between these values was generally irreconcilable. P10 mentioned being careful what they share on social media (*connection*) to stay safe (*security*). To improve social media, P10 also proposed decentralization of social media networks to create “close networks” based on one’s professional or personal links. P07 also wished to return social media such as Facebook back to a “true social media platform,” one that puts the value of *connectivity* above maintaining user attention.

P07: “Originally, like, in 2007, [Facebook] was just [to] connect with your friends. And [...] [posts] went into sequential order, and it was your friends, and there wasn't weird ads, and there wasn't like suggested reels [...] and like, it was just a true social media platform. [...] In 2007, Facebook was not the trash pile, I feel like [laughing] it is now. I just feel like Facebook is just like a literal, a literal dumpster fire. Like, I just I get on there [...] [just] to share pictures and keep my family informed of what I'm doing. Because that's what they use, like older people, [...] grandparents, aunts, and uncles [...] [and] it's just memes that no one's has researched. And it's just disinformation. And people's opinions that [...] don't even make sense. [Joint laughter] What is happening? [...] it's not people [being] like, “Oh, here's my family vacation.” [...] And I enjoyed that part of it, when it first came out, like people felt like they wanted to share about their lives, and share their own thoughts. And now it's just about sharing other people's thoughts. And like, having heated debates about current topics, like, that's what I feel like Facebook has become.” [Emphasis added]

Resolving this tension with “*double bind-like*” attributes by returning to a “true social media platform” would require reforming the attention economy. As described in Chapter 2, the attention economy fuels societal polarization using subversive tactics aimed at targeting our reactive selves over our higher cognitive capacities (Davenport & Beck, 2002; Goldhaber, 1997). It would also require us to move away from AI-fueled personalized social media feeds¹⁷⁸ – although doing so may require broader (possibly regulatory – or a mass public movement) interventions to accomplish this. Moving away from AI-curated feeds would also help us revolve another tension with *double bind* features linked to the attention economy – *Community (authenticity)* and *Use (time and convenience)*. Taken together, we can see that a viable alternative or solution to the attention economy will be needed in order to fully promote value-centered choices.

¹⁷⁷ P07, perhaps jokingly, also mentioned downloading mindfulness health apps like Calm to “balance out social media” as a strategy to tackle *connectivity* vs. *non-maleficence*.

¹⁷⁸ This does not mean that there can't be an aspect of personalization based on who you choose to follow – rather, the personalization is not done by AI. See, for example, Mastodon (<https://mastodon.social/explore>), where each “toot” by people you choose to follow are shown in chronological order.

c Minor Tensions

(1) *Minor Tensions: Overview*

There were six minor tensions which were mentioned in a more individualized manner, sometimes by as few as one participant. These were: *Control* vs. *Growth*; *Growth* vs. *Community*; *Growth* vs. *Use*; *Pleasure* vs. *Control*; *Safety* vs. *Growth*; and *Safety* vs. *Use* (“inter-value tensions”). In addition, there were three tensions within one value but between sub-values (“intra-value tensions”): *authenticity* vs. *connectivity* and *conformity* vs. *benevolence and universalism* within *Community*; and *time and convenience* vs. *utility and function* within *Use*.¹⁷⁹ For completeness, these are briefly presented here and suggested resolution strategies, where applicable, described. As we will see, there were instances where participants found it particularly challenging or were unable to resolve these tensions. These tensions had features of a *double bind* situation (“*double bind-like*”). The prevalence of *Control*, *Use*, and *Community* and their frequent tensions with mostly *double bind* attributes further suggest that these values are the *most* critical when thinking about promoting more value-centered privacy decisions.

(2) *Inter-Value Tensions*

Firstly, *Control* vs. *Growth*, discussed by four participants, focused primarily on the tension between meeting one’s health goals (*self-improvement*) using smartphone apps and the degree of control (*power and choice*) over their interactions with these apps. These participants found health apps, such as meditation apps, sleep trackers, period trackers, and fitness trackers helpful, but had some reservations. They firstly felt pushed to give more data or engage more with these apps due to gamified¹⁸⁰ app aspects. P14 expressed a tension between tracking to *self-improve* and the feeling manipulated by the app’s gamification tactics.

P14: “[App gamification] makes me uncomfortable [...] I know it's playing me [...] this is here specifically to engage me. It's not there for my benefit. It's there to make me use what I'm using [...] [while] it appeals to me, and I enjoy tracking certain things [...] [but] I also don't like this.” [Emphasis added]

The gamification aspects, such as flashy, persistent notifications (P07: “Every day come in and tell us” notifications) could also make some participants feel guilty for not engaging with the apps. Participants felt that “[they] should put that [data in the app] to see if it helps [...] [the app] can't really help [them] if [they are] not using it fully” (P07). This also applied to downloading apps, resulting in more data sharing than preferred. P29 described how their value of health and wellness could be used against them, to “fall for [the] trap” to download more meditation apps than needed and thereby share more data than they wanted to with the apps (Q27, Appendix IX). They go on to compare how this value “trap” is unique to the digital space, using the analogy of shopping.

¹⁷⁹ While these “intra-tensions” could suggest a lack of cohesiveness within a coded value, the sub-values still maintained the same state as their primary orientation. They were also outliers, sometimes only appearing in one interview.

¹⁸⁰ See footnote 42 for more on the ethics of gamification.

How Do We Value Data Privacy?

P29: “If, in my personal life, I had two of the same T-shirt. I wouldn't keep [it or], you know what, I wouldn't have [gotten] that [second shirt] in the first place. Because I'd go to the shop, and I would realize, “I already have this, I don't need it.” [...] it's even more important when [...] you're looking at apps, because like buying two T-shirts isn't going to affect my data privacy [...] [as much as] having the two apps. They're both taking information from me. I genuinely don't need them on my phone.”

Only one strategy to manage these tensions was put forth by P22, who aimed to use health apps with better data privacy policies (Garmin) than other apps (My Fitness Pal).

Growth vs. Community was mentioned by one participant (P08). For them, the tension was a conflict of *conformity (value: Community) vs. learning/staying informed (Value: Growth)*, where they expressed the pressure to use services such as YouTube and Google which came at the expense of the ability to think, to create, and to learn (Q28, Appendix IX).

Closely related to *Control vs. Growth* was *Control vs. Pleasure* because they both involve gamified apps. Besides the feeling of being manipulated to engage with or give data away to gamified apps that one is engaging with to *self-improve (Control vs. Growth)*, P14 also expressed the tension between apps gamification (“it hits”) to make them engage (*Pleasure*) and feeling like they were being pushed to share more data with the app than they would otherwise be comfortable with sharing (*Control*).

Growth vs. Use again involved health apps for meeting one's goals (*self-improvement*) and social media for *learning* new things (*Growth*). In this case, however, *Growth* was in tension with *time and convenience* and *utility and function*. For *self-improvement vs. utility and function*, two participants questioned whether health apps (P07, P29) are *truly* as helpful as they claim to be. For P07, this manifested as having to enter data into apps, such as a “period tracker,” becoming “too cumbersome [...] to use.” For P29, this manifested as downloading apps she heard recommended on podcasts or targeted ads, resulting in a bunch of similar health apps on her phone that may not be truly helping.

P29: “So it's like that you end up with this, like ecosystem of, of apps that correspond with like that value that I was talking about [...] your personal development and well-being. But actually, is it really doing anything for that? I'm not too sure, if you have five different apps that are all doing the same thing.”

Two participants additionally expressed that, while social media can be valuable for *learning* new things, it can also be *time-consuming* (Q29, Appendix IX). P08 describes this tension somewhat differently, pitting the *convenience* of Google against allowing us to *learn* new things (Q28, Appendix IX).

One participant found that health tracker apps caused them stress. This is the only instance of *Growth vs. Safety*, in this case, sub-values *self-improvement vs. non-maleficence*. For P28, there was a time when they tried health tracker apps (*self-improvement*) but then found “them a little bit oppressive [...] “this is your heartbeats,” or, you know, “you only had two hours of good sleep” or whatever.” For them, it was “just another thing to get stressed out [about],” harming them. P28 eventually resolved this tension by “steer[ing] clear” of downloading health apps, preferring to use “planner apps or the Pomodoro [app]” to meet their goals and *self-improve*.

Lastly, two participants exhibited tension between *Safety* and *Use*, albeit between different sub-values. The first of these participants, P08, described a tension between *time and convenience* and *non-maleficence*. While they enjoyed the *convenience* provided by apps and services such as Google Maps, they worried such convenience could be harming

How Do We Value Data Privacy?

us. They describe how they feel humans are made for a world “with friction” like “St. Andrews” in Scotland was made for golf (Appendix IX, Q39). In P08’s view, the real-world challenges us, and too much convenience could be causing us harm through “shrinking our hippocampus,” or reducing our ability to navigate and remember details ourselves.¹⁸¹ To combat this, P08 felt that, in an ideal world, algorithms could be designed to encourage us to engage with our “St. Andrews,” the real world around us. AIs could “incentivize people to get away from the screens for a while,” not just for “40 minutes,” but in a way that “gets them engaged in their own lives that’s immediately around them” perhaps using notices with more visceral images to nudge users offline. The second participant, P23, instead had a *Safety* and *Use* tension between *security* and *utility and function*. They frequently expressed their *security* concerns if asked to share their phone number (“that’s a bit dangerous”). Still, they felt that sharing their phone number was required to access certain services and features (“all of our activities nowa-days [relate] to the phone number”). Giving the example of Gmail, P23 notes that sharing their phone number can be helpful because if they “cannot access [their account], [Gmail] sends me [...] the link through the phone number.” In an ideal world, they would prefer to be able to access the service without sharing their phone number.

(3) *Intra-Value Tensions*

There were also a few instances of sub-values with the same, overarching value, in tension. While these share the central organizing concept, or directionality, of the overarching value, these intra-value tensions could be interpreted as a difference concerning which sub-value should be given most weight. There were three intra-value tensions, *authenticity* vs. *connection (Community)*, *conformity* vs. *benevolence/universalism (Community)*, and *utility and functionality* vs. *time and convenience (Use)*. Within *Community*, four participants described how the digital, online world increases *connectivity* sometimes at the expense of *authenticity*. For example, P06 described why they still decide to shop in person because of the bookseller’s recommendations.

P06: “I like the experience of going in [to the bookshop in] person. [...] **It’s just nice to go into a bookshop**, I get like to pick up some books, I get to have, like, have something in my hand and I get to talk to someone who works in the bookstore [...] [who] maybe can recommend me something. [...] I think people who go and learn how to become like a bookseller, or people who work in board game stores, I think they’re really good [at] what they do [...] I know when I go there, [...] **I’ll get a good recommendation from a person** and when I go back [...] I can go and be like, “oh yeah, you recommended me like x y, z book and I really liked that book.” And then there’ll be like, “oh, yeah, I also enjoyed it. And I also enjoyed this book because it was similar.” [...] Yeah, it’s just, **I know Amazon’s algorithm works the same way**. [...] it matches you with customers who have like, rated the same things high that you have rated, and then whatever they bought they recommended to you. [...] **but it’s not as personal because it’s just not**. I can’t really explain why it’s not as personal because it’s the same mechanism. [...] maybe it’s because there’s no person I can connect it to.” [Emphasis added]

In addition, P08 expressed concern that pressure to *conform* to the online world “is going to lessen our connection to nature and the world around us” (*benevolence and universalism*). For *Use*, one participant, P13, described liking apps such as Google Drive

¹⁸¹ The hippocampus is the part of our brain that plays a major role in memory and navigation.

How Do We Value Data Privacy?

for the *convenience* it provides them “to access [files] when [they] don't have the computer,” but are concerned about the drain it has on their phone’s battery (*utility and function*).

d Value Tensions Conclusions

Here, we have presented the values that seem to be the most contentious – with all the most prevalent tensions involving one or both of *Control*, *Use*, and/or *Community*. These tensions also appear frequently in a “*double bind*-like” manner, where no matter how the user or agent acts, they cannot act in full accordance with their personal identity (failure of *self-realization* and *self-unification*). Many of these binds have to do with a lack of relevant controls or alternative courses of action that fulfill both values. This suggests that resolving tensions between these most prevalent values could be the most critical issue for promoting value-centered choice. However, resolving the tensions that most resemble *double binds* may require broader or regulatory interventions aimed at modifying personalized news feeds and addressing the power of social media monopolies to fully allow for value-centered choices.

In addition, value tensions resolved along the *self-realization* or *self-unification* dimensions existed on a spectrum, with tensions most resembling *double binds* representing the most extreme violation of these dimensions. Participants also frequently mentioned ways to improve upon their resolution strategy of choice. We could, then, also explore supporting existing strategies (and designing for new ones) to facilitate the deliberative process of weighing and acting upon values. However, the concept of *double binds* – particularly relevant when considering value tensions between *Control*, *Use*, and *Community* – was helpful in identifying areas where users are *most* unable to act in full accordance with their values. These areas are the most problematic for making value-centered choices.

5.3.5 Value and Value Tension Interview Conclusions

Values are motivationally related to privacy decisions, driven by both an individual’s understanding of the value and the context in question. Even if the added complexity of a dimension-by-dimension analysis may not be fully necessary to capture the role of values in data privacy decisions, emphasis on *Control* (or lack of it), orientations towards commitments and motivations, and expected phenomena like the “apathetic user” are consistent with the 4DT-informed understanding of value-centered privacy decisions. In addition, tensions existed on a spectrum between resolvable tensions and “true” *double binds*. The most frequent values (*Use*, *Control*, and *Community*) were identified and appeared in frequent “*double bind*-like” tensions with each other. This suggests that these values are the most critical when considering designing for value-centered privacy decisions.

However, because many values were context-specific, this likely will not be sufficient. Perhaps the context specificity challenge could be captured by focusing on one tension: *Control* vs. *Use*. *Control* (especially sub-value *power and choice*) has been shown to be highly relevant across contexts, and *Use* (especially *utility and function*) is directly related to context (relevant sub-values). Perhaps, then, the value-centered approach could focus on initially resolving the major tension between *Control* and *Use*, which may help resolve (or at least reduce) other context-specific tensions. This, too, is explored more in the concluding chapter of this thesis as a (possible) future direction of this work.

While the emphasis of this investigation was on individual data privacy decisions, we can find, weaved in both values (*power and choice*) and value tensions (e.g., *Community and Use*), the need for broader intervention to support value-centered choices. As we saw, the interviewees had plenty to say about the systems we exist in – surveillance capitalism (Zuboff, 2019); attention economies (Davenport & Beck, 2002; Goldhaber, 1997); and social media platform dominance (Usman, 2022). These topics have already been touched upon concerning values and value tensions, such as the *Safety* of technology that is made to keep our attention; the feeling of being manipulated to give away data or by targeted ads (*Control*); the *power* held by certain companies and the need for decentralization; app gamified health apps that are “playing me” (P14; *Growth vs. Control*); doing away with personalized social media feeds focused on “bad news” that makes us feel like “the world is on fire” and return to a “true social media platform” (P07; *Community vs. Safety*); and feeling obligated to engage with dominant social media and messaging services, like WhatsApp (*Community vs. Control*). Companies, as previously mentioned when discussing *trust*, were also seen as doing things with data that the users did not consent to for profit making or sometimes “downright evil” (P01) purposes, such as Facebook selling data to Cambridge Analytica.¹⁸² Critically, tensions introduced in part or full by these structures were also the ones that *most* resembled *double binds*.

Other (likely regulatory) changes will therefore be needed to account for the current systems in which we exist, as well as the effects our individual data sharing can have on those around us. While it is out of scope of this work to provide a detailed analysis of privacy regulations, it is important to consider the limits of the value-centered approach within our current political and economic systems, especially considering the concerns raised by interview participants. In addition, while there are genuine concerns about the effectiveness and desirability of our current privacy regulations that are already explored at length in existing literature,¹⁸³ we must still aim to account for the broader implications of an individual’s data privacy decision. Their decision can also impact our a societies and our democracies (Cohen, 2013; Zuboff, 2019), and accounting for this broader dimension of privacy decision-making harms will likely require regulatory intervention. Broader regulatory interventions, and their interplay with systems targeted at the individual, such as the VcPA built with the value-centered approach, will therefore be a critical area of future research. This will be explored in greater detail in Chapter 7.

Section 5.4 Conclusion

Based on the online survey (Phase I) and the semi-structured interviews (Phase III), we can now return to **RQ1**: *What is the relationship between values and privacy preferences when deciding to download an app, if any?* (Table 4-1 in Chapter 4). In particular, we wished to answer *how we value privacy* in order to: 1.) better understand how values are involved in privacy decision-making as a means of promoting more value-centered choice; and 2.) to evaluate the 4DT-informed value-centered approach.

To start with (2), the interview data supports the 4DT-informed understanding of values and data privacy decisions – even if the added complexity of a dimension-by-dimension analysis may not be fully necessary to capture the role of values in data privacy decisions. Emphasis on *Control* (or lack of it), orientations towards commitments and

¹⁸² See footnote 144 for more on the Cambridge Analytica scandal.

¹⁸³ For example, initial empirical data suggests that smartphone data tracking on Android devices has not significantly changed following the introduction of the GDPR (Kollnig et al., 2021), greatly drawing into question its effectiveness. Another powerful study found that *none* of the 400 popular apps they studied tracked users in accordance with EU privacy laws (Paci et al., 2023).

How Do We Value Data Privacy?

motivations, and expected phenomena like the “apathetic user” are consistent with the 4DT-informed understanding of value-centered privacy decisions, supporting our use of the theory to identify design features of a privacy assistant. We also found that value tensions existed on a spectrum between resolvable tensions and “true” *double binds*. In particular, the presence of structural factors – such as social media monopolies, the attention economy, and surveillance capitalism – were defining features of tensions that *most* resembled *double binds* (those that were most “*double bind-like*.”) In addition, the interview data suggests that power structures are related to the value of *power and choice* by causing *self-realization*, *self-unification*, and *self-constitution* failures (in the case of the “apathetic user”). We also found that the *inertia bias* can cause users to stick with old apps on their phone (failure of *self-realization* and *self-unification*). These additional insights can be accommodated by 4DT, added to our initial conceptualization from Chapter 3. However, fully addressing challenges introduced by structural factors will likely require that the value-centered approach here is complemented by a broader approach.

Considering the results of Phase I and the (relevant) results from Phase III, while general themes became apparent, overall values are involved in privacy decisions and apps choice in a highly individualized, context-specific manner. We observed that different values were more relevant based on the app in question in Phase I (e.g., *benevolence and universalism* for OpenLitterMap) and in the interviews in Phase III (e.g., *Community* and sharing data with/engaging with social media). Value relationships to privacy decision-making, both more traditional notice-and-consent privacy decisions and choosing a smartphone app, were highly individually variable in both the survey and the interviews. This suggests that the value-privacy relationship is highly informed by individual preferences and understandings of values. As discussed, this complex state of affairs may not be best captured in a survey based on the very general values of the TBHV.

These individual differences and method limitations further suggest that measuring and understanding individual value-privacy relationships well enough to promote value-centered choice will be very challenging. However, focusing on the most prominent values (and their stronger tensions with *double bind* features) identified in the interviews – including *Control* (especially *power and choice*), *Use* (especially *utility and function*), and *Community* (especially *connection*, mostly social media and messaging services) could be one possible way forward. Initially focusing on resolving the major, frequent tension between *Control* and *Use* could help resolve (or, at least reduce) important context-specific tensions due to *Use*'s close association with context-specific values.

All of these insights also have relevance for designing a value-centered privacy assistant (VcPA). In the next chapter, we will explore how our prototype VcPA system was received by users (**RQ2**) and start to weave **RQ1** and **RQ2** insights together to inform future VcPA design.

Chapter 6 Evaluating the Value-Centered Privacy Assistant

*Try not to mistake what you
Have with what you hate
It could leave, it could leave
Come the morning celebrate the night
It's the fall before the climb
Shall we sing, shall we sing
'til the morning?
[...]
C'mon, c'mon
With everything falling down around me
I'd like to believe in all the possibilities*

Fun and Panic! At the Disco (“C’Mon ”)

Section 6.1 Chapter Overview

This chapter explores the results of the Mock App Store (Phase II) and follow-up semi-structured interviews as proof-of-concept that a value-centered privacy assistant (VcPA), designed using privacy preferences *and* values, could help users make privacy decisions such as choosing apps (RQ2: Table 4-1, Chapter 4). We evaluate the desirability and effectiveness of the value-centered privacy assistant (VcPA) at helping users make app choices more consistent with their values, as well as elicit feedback on the prototype,¹⁸⁴ in order to lay the groundwork for future VcPAs. Participants engaged with the VcPA, consisting of selective notices, a “suggest alternatives” feature, and exploratory notices (Table 3-2, Chapter 3) in a synthetic, online app store called the Mock App Store (MAS). We establish that a value-centered approach to privacy decision-making, operationalized as VcPA, can help serve as a form of *self-binding* by helping users make more value-centered app choices. A particularly well-received feature was “suggest alternatives,” which some participants found helped them find alternative apps of similar function but more consistent with their values. While this study supports proof-of-concept – that is, that value-centered privacy assistants can be helpful to users by promoting value-centered decisions – we identified three points of improvement. These were: VcPA profile creation, a more streamlined “suggest alternatives” feature, and a more user-friendly privacy notice presentation. Our results indicate future research could focus on constructing more personalized profiles; building profiles based on (non-Schwartz) values; using other (non-Apple Privacy Label) privacy ontologies and representations on VcPA notices; and a “suggest alternatives” page embedded into the selective notice itself. Longitudinal studies with VcPA users will also be needed to further tune the timing of exploratory notices, thereby ensuring “exploration” that counters the *inertia bias* without causing notice fatigue. Future research into these areas will be critical for moving the VcPA from proof-of-concept into an assistant that can be deployed on smartphone app stores.

¹⁸⁴ The prototype was constructed using the survey data from Phase I, as described in the Chapter 4.

6.1.1 Collaborator Contributions

The studies presented in this chapter were conducted in collaboration with Prof. Dr. Mathieu d'Aquin (supervision guidance, Mock App Store implementation, value-centered privacy assistant implementation, data analysis, manuscript feedback), Dr. Heike Felzmann (supervision guidance and manuscript feedback), Prof. Dr. Kathryn Cormican (supervision guidance and manuscript feedback), Dr. Dave Lewis (supervision guidance), Dr. Ilaria Tiddi (supervision guidance and Mock App Store implementation), and Dr. Dayana Spagnuolo (data analysis and value-centered privacy assistant implementation). I (the PhD candidate) conducted the study, as well as worked with collaborators at all stages of data collection, data analysis, and results write-up.

6.1.2 Relevant Papers and Conference Contributions

Some material in this chapter, including certain text and figures, has been previously published or presented in the following:

Carter, S.E., d'Aquin, M., Spagnuolo, D., Tiddi, I., Felzmann, H., & Cormican K. (2023) The privacy-value-app relationship and the value-centered privacy assistant. ArXiv. <https://arxiv.org/abs/2308.05700>

Carter, Sarah E., Tiddi, Ilaria, & Spagnuolo, Dayana. (2022, June 13). A “Mock App Store” interface for virtual privacy assistants. Hybrid Human Intelligence 2022: Augmenting Human Intellect (HHAI2022), Amsterdam, the Netherlands. Zenodo. <https://doi.org/10.5281/zenodo.8204393>

Carter, S. E., Tiddi, I., & Spagnuolo, D. (2022). A “Mock App Store” interface for virtual privacy assistants. In S. Schlobach, M. Pérez-Ortiz, & M. Tielman (Eds.), *HHAI2022: Augmenting Human Intellect* (Vol. 354). IOS Press. <https://doi.org/10.3233/FAIA220212>

6.1.3 Participant Demographics

Following the Phase I survey, a second group of participants were recruited to partake in the Mock App Store Study (Phase II). For the Mock App Store Study, we obtained 120 engagements. Of these engagements, 111 participants completed the entry survey with demographic details (Appendix VI). Participants were primarily adults (ages 25-64, 82 participants). 25 participants were young adults (18-24); 2 were older adults (65+), and 2 preferred not to say. Roughly half (63) of participants identified as women, with 42 identifying as men and 6 as other/non-binary/prefer not to say. Nationalities were grouped by continent, the majority of European nationalities (67), followed by Asia (21), North America (13), and 10 other/prefer not to say. Fluent and native English speakers were split evenly (49 and 62, respectively), and all participants currently owned a smartphone. The majority also had (or were in the process of obtaining) a doctoral or master's degree (81), with 19 for bachelor's degree, 7 for a secondary degree, and 4 preferring not to say. After excluding logs from the MAS that did not include any downloaded apps, we had the logs

of 77 participants who completed the Mock App Store exercise. 66 participants completed the exit survey after the exercise.¹⁸⁵

Relevant interview results are also discussed in this chapter, and the demographics for the interviews are described in Chapter 5.

Section 6.2 The VcPA and Value-Centered App Choices

Overall, the VcPA appeared to help participants download apps more consistently with their values. In a VcPA, selective notices are issued to users when an app's data collection practices are inconsistent with the user's values (as captured by their selected profile). To operationalize this, recall from Chapter 4 (Methods) that an acceptability coefficient for each profile-app pair was calculated to determine when selective notices would be presented, with a cutoff set to <0.1 .¹⁸⁶ We can therefore calculate how many apps "downloaded" on the Mock App Store (MAS) have a coefficient >0.1 with the user's selected profile as a means of determining how many apps were downloaded consistently with a user's values. In this case, 35 of participants had a high percentage ($>90\%$) of apps downloaded at the end of the Mock App Store exercise match their profile, with only 11 having a low percentage ($<10\%$ match) (Figure 6-1). Interestingly, those who selected the Helpful Neighbor profile tended to download more apps that were consistent with their profile than the other two profiles ($p=0.0003$ and 0.0008).¹⁸⁷ This suggests that the VcPA's selective notices were reasonably successful at acting as a form of *self-binding* – helping participants to act according to their values (uphold *external self-realization* and *self-unification*).

Section 6.3 VcPA Profiles

Participants leaned towards finding values associated with each profile clear and finding a profile that reflected them (Figure 6-2). In the interviews, three participants felt they found a profile that was a good match for them (P09: "I suppose really, like Goal Setter was very, very, very close to perfect, probably perfect"). Profiles were also interpreted similarly. Goal Setter was (positively and negatively) seen as primarily focused on *work* goals (5 participants; P14: "Yeah, I'm more of a Not Goal Setter!") rather than goals more broadly as was intended (3 participants). In this vein, some participants saw Goal Setter as being about the destination rather than the journey (P22: "I enjoy the process of doing things as well. But in terms of setting goals, I just didn't like it that it wasn't really talking about, you know, the journey towards getting there") or about living in the future, rather than the moment (P08: "I find myself a bit more out on the limb most of the time, and [...] [that's]

¹⁸⁵ Unfortunately, there was a large drop-off of participants between completing the first stage of the Mock App Store Study – an entry survey – and starting engagement with the store (Chapter 4). In the interviews, three participants reported never making it to the Store after completing the entry survey (Q38, Appendix IX). The link and directions at the end of the survey should have probably been clearer, perhaps bolded, to catch participants' attention. There was also a (smaller) drop of participants from the Mock App Store to the Exit survey (77 to 66), perhaps because they also failed to click on the link on the final "thank you" page.

¹⁸⁶ All acceptability coefficients are available in Appendix II.

¹⁸⁷ A 2 component PCA analysis was also conducted between profiles and the following features: the number of downloads; the number of apps added to their virtual smartphone; the number of apps deleted from their virtual smartphone; the number of selective notices received; the number of exploratory notices received, the number of times setting a profile; the number of apps downloaded outside their profile; the % of apps at the end of the study in their profile; and the percent of times "see alternatives" was used when receiving a selective notice. No other patterns, however, were observed between the profiles and these features. For this analysis, those participants that changed profiles partway through the study (4 participants) were ignored.

where I enjoy being.”). Demographic considerations, such (younger) age or (higher) education level, were also occasionally associated with Goal Setter. The values *Hedonism* and *Power* evoked negative associations and motivated participants to not pick Goal Setter, even if they related to the value of *Achievement*. Adventurer was primarily associated with individual autonomy and freedom to live one’s life (4 participants), as well as being open to living new experiences (3 participants). For example, for one participant, it was associated with being easy-going, or “just go[ing] with it” (P12). For some, how young one was as well as their occupation (e.g., student) was positively associated with Adventurer. Helpful Neighbor was viewed (both positively and negatively) as being associated with being conforming, humble, and caring for others first and foremost. This suggests that the three VcPA profiles we developed were generally perceived as understandable and contained many of the values that mattered to users.




Figure 6-1: Number of participants who downloaded profile-matching apps x - y% (in decimal) of the time¹⁸⁸

6.3.1 VcPA Profile Improvements

While the VcPA helped some users choose more value-consistent apps and profiles were generally well-received, our study also provides insights into how they could be improved. In particular, more personalized profiles and improving the accuracy of the value-privacy relationship on which they were based could be further optimized. These identified improvements also point us to areas for future research into the VcPA.

a More Personalized VcPA Profiles

While participants generally reported finding a profile that matched them, we also learned that many interview participants saw themselves between profiles. While it is notable that two participants interpreted profiles as fundamentally in opposition with each other,¹⁸⁹ most participants felt there were significant areas of overlap: 4 were between Helpful Neighbor and Goal Setter; 2 between Adventurer and Helpful Neighbor; 1 between Adventurer and Goal Setter); and 4 between every profile. Three participants described the profiles as “high-level” (P02) and vague, with the line between profiles not clear. Two

¹⁸⁸ E.g., 35 participants downloaded apps consistent with their profile 90% - 100% of the time.

¹⁸⁹ One participant viewed Helpful Neighbor’s *Security* as fundamentally opposed to Adventurer and another viewed Helpful Neighbor’s emphasis on humility as fundamentally opposed to Goal Setter’s focus on success.

interview participants further felt that the profiles and notices generated from them were not reflecting their value and privacy preferences, and one felt that their real-life profile would be a mix of profiles. Participants utilized various strategies to try to choose the best profile for them. For some, the profile was selected based on the best “all considered” choice or a process of elimination, while for others, it was picked based on their age, their gut feeling, or even their mood when taking the survey (Q34, Appendix IX). One participant also felt that the order in which the profiles were presented may have influenced their choice when seeing themselves in multiple profiles, picking the one that was presented first (Adventurer).

Taken together, these results suggest that VcPA profiles could be improved by being more personalized. Our initial profiles did help users make more value-centered choices, serving their role in establishing proof-of-concept of a value-centered approach. Designing more personalized profiles, however, could further help users act according to their values, moving us from proof-of-concept to a more optimized VcPA.

To guide us forward with this improvement, we also identified some participant-proposed improvements. These included: customizable profiles to allow users to create the profile that best reflects them, a survey to sort a user into a profile, and adding more profiles. To avoid cognitive overload with an excessive number of profiles, we think that designing customizable profiles (“mix-and-match”) and/or using a survey to initially sort users into profiles would be the most promising avenues of future work. This will be an important area of future work and possible avenues to accomplish this are explored in greater detail in Chapter 7 (Conclusions and Future Directions).

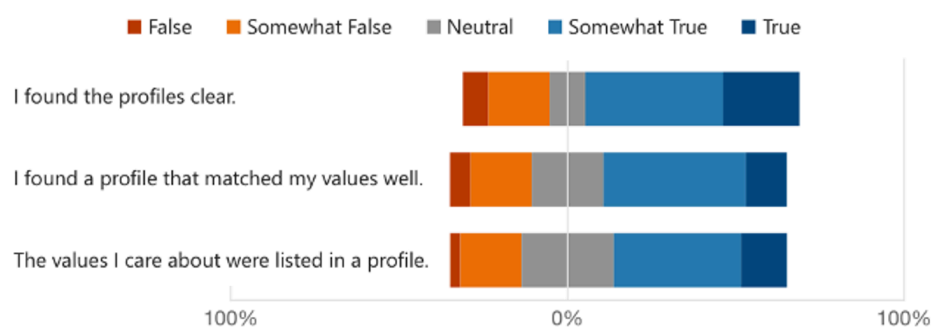


Figure 6-2: VcPA profile reception

b Basing Profiles on Control, Use, and Community

The value-privacy relationship on which the profiles were based does not appear to be capturing *all* relevant values, and the relationship between values and privacy preferences was not always clear to participants. This suggests that profiles could be designed using another method than one based on the TBHV – one that is clearer and better captures the values relevant to app choice and privacy decisions.

Firstly, the relevance of other values was present in the data in a few ways. To start, it was indicated by the comments left by those who ignored notices, who indicated that other considerations (implicitly or explicitly informed by values) outweighed privacy or other concerns that they had.¹⁹⁰ For example, one MAS participant wrote: “The app is a brain trainer, so even if it has trackers, it helps in stimulating my senses creating a sense of adventure” (Appendix XI). This suggests that this participant’s value of *stimulation* was

¹⁹⁰ For a deeper discussion of study results related to VcPA notices, see Section 6.4.

not being accurately captured by their profile (even though it is part of Adventurer). Two interview participants further felt that the profiles and notices generated from them were not reflecting their values and privacy preferences. As we saw when answering **RQ1** (in the survey and interview data), participants also felt that their relevant values vary based on the context. Four interview participants felt like their relevant values would be heavily dependent on context (e.g., privacy or choosing an app) rather than based on global, life-guiding principles. One felt that the link between privacy preferences and values in the profiles was perhaps not there at all (P06) (Q35, Appendix IX). Another felt that the values under the profile did not match the profile's title or description. One actively resisted the link between profiles and data privacy, because "who are you to tell me what my values are?" (P04). How participants described searching for apps also suggests that some relevant values were not being captured. Participants describe other considerations that are intrinsically value-laden but, in this case, were considered distinct from the value profiles. In particular, app *usefulness* and *function* were a critical consideration when choosing apps in the app store for nearly half (7) of participants. Some participants took a "function first, privacy second"¹⁹¹ approach to searching for apps (Q36, Appendix IX), similar to the approach used when searching for apps in real life (*Use vs. Control*, Chapter 5). Because of the exercise prompt and content of the MAS, the most stated desired function was to help meet health goals. In addition, participants also stated choosing apps according to what they value (2 participants) or to minimize data collected based on their preferences (4 participants). These values and privacy preferences, however, were sometimes considered separately from their profile, suggesting the profile was not always serving its role to help them make value-consistent choices. Others found apps attractive based on having an "appealing name" (P30) while two others valued choosing apps in a more intuitive, quicker process, suggesting that app *attractiveness* and *efficiency* of choosing an app were not fully captured by their VcPA profile. P14 further suggested adding the value *engagement*, which they find especially relevant when choosing apps. Taken together, these results indicate that there are other likely context-dependent values that were not being captured by the more general profiles.

Secondly, lack of clarity around the value-privacy preference relationship was indicated by decreased participant trust that their profiles were helping them make app choices based on their values and observed profile misunderstandings. When asked how confident the participants were that they made app choices consistently with their values on a scale of 1-5, only 56% of participants who completed the exit survey rated >4 (5 rated 1; 13 rated 2; 13 rated 3; 25 rated 4; 10 rated 5). This is despite that, as reported previously, the VcPA did succeed in promoting downloading apps that were consistent with participants' profiles. We further learned that there were some misunderstandings about the profiles (even though participants largely reported that they found the profiles understandable). It was unclear whether the profiles were meant to be values in the specific context of privacy (5 participants) or intended life-guiding principles (5 participants). In addition, it was unclear to some participants if they should pick a profile based on how they see themselves, the "ideal" version of themselves, or how others see them.¹⁹² Based on our goal of promoting value-centered privacy decisions, the intention of the profiles was what is ideal to the participant (but what might be currently unfulfilled when making everyday data privacy or app choices). Some also misinterpreted the relationship between

¹⁹¹ Four participants also described this ranking in terms of whether they decided to choose an alternative app with better profile consistency from the "select alternatives" page. This is discussed more in Section 6.4.

¹⁹² Profiles were also once misinterpreted as personality profiles rather than value profiles, with one participant (P30) choosing the Adventurer profile because it was the closest match to their Myers-Briggs (<https://www.myersbriggs.org>) result.

values and data privacy on which the profiles were based. It was seen by a small handful of participants as being more privacy-preserving to protect oneself (*security*) and others' privacy, while another felt Adventurer meant being more "adventurous" with data-sharing. This was not the case – every profile had apps that collected data and were still considered consistent with the profile.¹⁹³

Based on these results and the feedback received, future VcPA profiles will need to consider the context-specificity of values. In addition, the value-privacy preference relationship on which they are based will need to be made clearer. Likely, these challenges are linked to the Schwartz survey limitations used to build the current profiles, which, as we found in the survey results in Chapter 5, struggle to capture the context-dependent nature of values in privacy decision-making. Better profile design will therefore require a deeper understanding of the value-privacy-app relationship outside the TBVH. We started this exploration through the interview investigation into how values are involved in privacy preferences in Chapter 5, finding that the values *Control*, *Community*, and *Use* were quite prevalent in the interview data. We further propose that, due to *Use*'s close association with the app context (and more context-specific values), that resolving the tension between *Control* vs. *Use* may be one avenue of accounting for context-specificity. In this vein, profiles could be designed based on *Control*, *Use*, and *Community*. More context-specific methods for value identification, such as using AI and human annotators to identify context-specific values (Liscio et al., 2022), could be another possible path forward to identify relevant values for profile design. To improve the clarity of the value-privacy preference relationship on which the profile is based, we could present a clear mapping of the privacy preference to the value on the privacy notice itself (to be touched on more in the next section, which looks at VcPA notices). It is also possible that moving away from difficult-to-understand Schwartz values will also increase understandability. Each of these possible areas of future direction to improve profile clarity and the value-privacy preference relationship will be explored in greater detail in Chapter 7 (Conclusions and Future Directions).

Section 6.4 Selective Notices and the "Suggest Alternatives" Feature

Selective notices and the "suggest alternatives" feature were well-received by participants, showing a proof-of-concept that selective friction and offering an alternative course of action can help users act in better accordance with their values.

Firstly, the "suggest alternatives" feature was the most popular feature, with 69% of Mock App Store participants rating it 4 or higher (average=3.79, SD=1.28, median=4) (Table 6-1). In addition, a few interview participants expressed satisfaction with the "suggest alternatives" page, describing it as "really handy" to consider "apps with the same functionality" (P14), something they "would never think to do [...] on [their] own" (P07). It therefore appears that the "suggest alternatives" feature fulfilled its function of offering an alternative course of action for users to act upon (and thereby *self-realize* and *self-unify*).

¹⁹³ Based on the acceptability coefficient of <0.1 (see Appendix II for a list of all the acceptability coefficients)

Table 6-1: VcPA feature reception (Likert scale, 1-5)

Question	1	2	3	4	5	Average	SD	Median
In general, how helpful did you find the notifications that you received?	2	11	14	21	18	3.64	1.15	4
How helpful were the timing of pop-up notifications?	1	6	13	19	27	3.98	1.06	4
How did you find the frequency of the notifications?	2	5	41	12	6	3.23	0.84	3
How helpful did you find the "show me alternative applications" feature?	6	6	8	22	24	3.79	1.28	4

Secondly, most aspects (timing, content, frequency, overall) of the VcPA notices were moderately positively received (Table 6-1). For the notices, 80% of participants rated their overall satisfaction with VcPA notices above an average score of 3 out of 5 (average =3.64; SD=1.15; median=4). 62% rated the frequency of notices above 3 (average=3.23, SD=0.84, median=3). 89% rated the timing of the notices above 3 (average=3.98, SD=1.06, median=4).

It is also notable that the questions asking participants to rank their level of overall privacy concern and smartphone privacy concerns, included in both the entry and the exit survey, were not significantly different ($p=0.83$ and 0.37 , respectively). This suggests that feedback on the exit survey about the Mock App Store was not due to a task-stimulated change in privacy concern. However, those who ranked their privacy concern higher (in general or on their smartphone) on the exit survey were significantly correlated with higher ratings for all VcPA modalities except timing (correlation range: 0.27 to 0.38). This suggests that those who are more concerned about their privacy tend to be more satisfied with VcPA features. This makes some sense if we recall the “privacy concerned” user scenario from Chapter 3, where the VcPA selective notices best assisted this user.

a Streamlined “Suggest Alternatives” Feature

While the “suggest alternatives” feature appeared to help participants find alternative apps, we also learned that there were a few places where it could be improved in future work. When asked what other features would help them select a smartphone application, some Mock App Store participants recommended improving the ability to compare apps on the “suggest alternatives” page with added information about the app (e.g., app reviews) and app functionality (Appendix X). Some participants also did not think the alternatives recommended were close enough in terms of their function. Feedback gathered from the interviews echoed that in the written feedback, including difficulty comparing alternatives and a lack of function match between supposed similar apps (Q37, Appendix IX). In addition, one interview participant mentioned that it would be helpful to include a side-by-

side comparison of paid vs. free features, and another participant recommended that the notices be used to provide greater transparency around how apps are vetted by an app store.

The “suggest alternatives” feature would therefore benefit from more app information; more ease in comparing similar apps; and better matching the similar app functions to the original app. The first limitation concerning app information was more a limitation of the MAS interface, something that would likely not be an issue if a VcPA was used in a “real” app store setting. However, a comparison feature, perhaps on the selective notice itself, could further enable users to make app choices consistent with their values. Matching based on descriptions, rather than keywords, could also improve the function match, and could be a promising area of future research. We will return to these avenues of future research in Chapter 7.

b Clear Presentation of Privacy Preferences

While two interview participants reported that the notices caused them to pause and reflect more on data privacy¹⁹⁴ and notices were well-rated, the use of the “suggest alternatives” page was split between two extremes. Roughly a fourth of the participants had high engagement (using the “suggest alternatives” button >90% of the time they received a selective notice), while a fourth of participants had low engagement (clicking less than <10% of the time). This suggests that the “see alternatives” button on the selective notice was especially engaging for some participants but equally disengaging for others. However, its high rating by MAS participants suggests that the feature was a desirable one.

To see how notice engagement could be improved, we explored why notices were ignored by some participants. When asked on the Mock App Store exit survey why participants ignored the selective notice and decided to download an app anyway, most respondents felt that the utility or function of the app to obtain their goals was more important than data privacy concerns (Appendix XI) (example: “The app is a brain trainer, so even if it has trackers, it helps in stimulating my senses creating a sense of adventure.”) Notably, one participant did not think that their choice “means much” (“I don't feel the relationship between the value profile and my choice means much. The app has utility I'd like, the privacy concerns are small”) and another expressed confusion about the type of data being collected (“I am not sure what tracked, linked, unlinked means, so I am not sure what they collect exactly.”) Written feedback also contained a request for greater clarity on what data is being collected by an app. These sentiments were further echoed in the interviews, where one participant stated that it was not clear what those data types mean. Another believed that linked and tracking data were still anonymous, and the lack of clarity about what unlinked/linked/and tracking means were also expressed as a reason for ignoring selective notices (Appendix XI). Interview participants also mentioned that they ignored notices because the explanation (link between privacy preferences and values) on the notice was unclear or too long; the function of the app was highly desirable; or simply because they were annoying (P28: “swatting a fly”). The notices were also sometimes misunderstood to be about data collection only rather than about values. We additionally saw misunderstandings of the “traffic light” feature that gave them a visual representation of how consistent the app is with their profile (see Chapter 4). The stoplight was misinterpreted by two participants as less data collected instead of greater consistency and another participant thought it could be more intuitive if you could hover on the stoplight to

¹⁹⁴ This reflection caused one of them to change their behavior (not download the app), while the other decided to continue to download the app anyway.

get more information about what data is being collected. One participant found the stoplight more helpful than the pop-up notices (“big blurb of information”) because it was a “short, really concise bit of information that allows you to just see, in clear terms, what’s being taken” (P29).

While some of these results are likely tied to profile improvements around the values, our results suggest that the Apple Privacy Label ontology we utilized to represent privacy preferences is difficult for users to understand. This supports other reports that the labels are not as effective as once hoped (Kollnig et al., 2022; Zhang et al., 2022).¹⁹⁵ The confusion could also have contributed to the weak correlations from the survey (Chapter 5), skewing the results if participants interpreted the privacy preferences differently. While it remains a challenge to hold app developers accountable for providing correct information (Ali et al., 2023; Jain et al., 2023; Koch et al., 2022; Kollnig et al., 2022; Rodriguez et al., 2023),¹⁹⁶ improving the presentation of privacy preferences on VcPA notices will be critical for understandability and transparency. We will therefore need to explore other privacy ontologies and representations. While there are many considerations when it comes to making data collection explanations more understandable,¹⁹⁷ for our purposes, we could perhaps utilize an approach similar to the “Privacy Facts” display described in Kelley (2013). Their approach includes a subset of privacy details as well as other relevant information for apps, such as reviews, on their notice. As mentioned previously when looking at profiles, we could visually display the data being collected alongside the associated values to increase user comprehension. A greater investigation of possible privacy notice presentations and future directions is provided in Chapter 7.

Section 6.5 Exploratory Notices

Lastly, it was exceedingly difficult to test the exploratory notices. Recall that the goal of exploratory notices is to account for the *inertia bias* – that is, instances when one’s values have shifted and the profile they have selected is no longer a good match for them. Recall also that tuning the timing of these notices is critical to balancing their intended function with inadvertently contributing to notice fatigue. Given that value changes are something that generally happen over a longer period of time (and, most certainly, quite infrequently within the five-minute timeframe of the Mock App Store Study), we were not able to gather much insight into their effectiveness. Indeed, no one interviewed seemed to recall getting one, and no feedback left on the exit survey specifically mentioned them. However, the overall moderately positive reception of notice timing (Table 6-1) suggests that having one triggered during a 30 second interval of a five-minute time period was not *particularly* bothersome (see Chapter 4). Longitudinal studies with VcPA users will be a critical area of future research to further tune the timing of exploratory notices, balancing the “exploration” that counters the *inertia bias* with excessive friction that may encourage a value-inconsistent choice.

¹⁹⁵ It is also notable that one participant – P23 – found the difference between unlinked, linked, and tracking data unclear and confusing when completing the *survey* (Phase I). This, in addition with the confusion around the survey value questions (Section 5.2) could have contributed to the high variability of survey results. It further supports the conclusion drawn here from the MAS (Phase III) results that the privacy label ontology is not the most understandable.

¹⁹⁶ Significantly, Jain et al. (2023) recently reported that 88% of the 354,725 Apple apps they surveyed had at least one discrepancy between its privacy label and privacy policy.

¹⁹⁷ For a summary of these different considerations, see Schaub et al. (2015).

Section 6.6 Conclusion

RQ2 aimed to establish the usability and effectiveness of the value-centered privacy assistant and to identify areas of future work to improve the VcPA. The results here accomplish the proof-of-concept that a value-centered privacy assistant, designed using privacy preferences *and* values, could help users when making privacy decisions such as choosing apps by allowing them to *self-bind* and act (*self-realize*, *self-unify*) according to their values. A particularly well-received feature was “suggest alternatives,” which some participants found helped them find alternative apps of similar function but more consistent with their values. Overall, participants also found profiles clear and could find one that they felt reflected them. Selective notices were also moderately positively received. In addition, we saw that the VcPA did by-and-large help participants download more apps consistent with their values (in this case, defined by an acceptability coefficient for the profile-app pair that is greater than 0.1).

We also succeeded in identifying areas where the VcPA could be improved – laying the foundation for future research into VcPAs. Future research could explore how VcPA profiles could be more tailored to each participant, perhaps by making them customizable and/or using a survey to sort users into profiles, as was recommended by some of our study participants. The value-privacy preference relationship on which the profiles are based will therefore require investigation outside the Schwartz understanding of values and instead encompass value context-specificity. Building profiles based on the most prevalent *Control*, *Use*, and *Community* values identified as important to users in our interview investigation (Chapter 5) or looking at utilizing more context-specific methodologies for exploring values in technology could be two potential avenues of investigation. Future research could also investigate designing a more streamlined “suggest alternatives” feature on the notice itself with better function match, as we found was desired by participants. Clearer presentation of values and data collection practices on selective notices, perhaps by visually displaying the data being collected alongside the associated values, could be explored to increase user comprehension. And lastly, longitudinal studies will be needed to further fine-tune the timing of the exploratory notices to balance accounting for the *inertia bias* against notice fatigue concerns.

Having further understood how we value data privacy (Chapter 5), established proof-of-concept for the VcPA and the value-centered approach to data privacy decisions, and identified avenues of improvement based on user feedback on the VcPA, I will now conclude by outlining the significance of this work in progressing a value-centered approach to data privacy decision-making and present possible avenues for future work in greater detail.

Chapter 7 Conclusion and Future Directions

*Come writers and critics who prophesize with your pen
And keep your eyes wide, the chance won't come again
And don't speak too soon, for the wheel's still in spin
And there's no tellin' who that it's namin'
For the loser now will be later to win
For the times, they are a-changin'*

Bob Dylan (“The Times They Are a Changin’ ”)

Section 7.1 Chapter Overview

In this final chapter, I conclude and present future avenues for the value-centered approach to data privacy decision-making. I start with an overview of the major contributions of this thesis – conceptualizing value-centered privacy decisions; designing a smartphone value-centered privacy assistant (VcPA) prototype; and empirically interrogating this value-centered approach using a mixed-methods investigation. Then, I dive deeper into the empirical insights to propose areas of future research into value-centered privacy approach, including notice and profile design; identifying relevant values; harmonizing an individual, value-centered approach with broader privacy approaches such as regulation; and designing VcPAs for other privacy contexts.

7.1.1 Collaborator Contributions

The ideas described in this chapter are my (the PhD candidate’s) work. Feedback was provided by PhD supervisors Dr. Heike Felzmann, Prof. Dr. Mathieu d’Aquin, Prof. Dr. Kathryn Cormican, and Dr. Dave Lewis. Collaboration with and feedback from Xengie Doan (University of Luxembourg) and Marcu Florea (University of Groningen) were helpful in initially exploring and applying the value-centered approach to other contexts.

7.1.2 Relevant Papers and Conference Contributions

Some material in this chapter, including certain text and figures, has been previously published or presented in the following:

- Doan, X., Florea, M., & Carter, S. E. (2023). Legal-Ethical challenges and technological solutions to e-health data consent in the EU. In P. Lukowicz, S. Mayer, J. Koch, J. Shawe-Taylor, & I. Tididi (Eds.), *HHAI 2023: Augmenting Human Intellect* (pp. 243–253). IOS Press. <https://doi.org/10.3233/FAIA230088>
- Carter, S.E., d’Aquin, M., Spagnuolo, D., Tididi, I., Felzmann, H., & Cormican K. (2023). The privacy-value-app relationship and the value-centered privacy assistant. ArXiv. <https://arxiv.org/abs/2308.05700>

Section 7.2 Overview of Research Findings and Contributions

When faced with so many privacy decisions, we often struggle to make meaningful privacy decisions. This is especially difficult when privacy notices are being designed in manners that exploit our cognitive biases and heuristics, coaxing us to consent to data sharing. Much research has been done into privacy self-management and the use of these design tricks, calling into question the effectiveness of privacy notice-and-consent regimes at eliciting informed consent. This work aimed to return to the normative basis of informed consent – to respect autonomy. Instead of aiming to elicit informed consent, this work aims to understand respecting autonomy in data privacy decisions as promoting value-centered privacy choices – that is, choices centered on our personal values. To this end, this thesis defined, conceptualized, interrogated, and designed for value-centered privacy decision-making as a means of respecting and promoting autonomy. This work lays the groundwork – from theory to practice – for future computer science researchers to design for value-centered privacy decisions.

In **Chapter 3**, I first **conceptualized and defined value-centered privacy decision-making** using a value-centered theory of autonomy – the Four-Dimensional Theory of Self-Governance (4DT) (Table 7-1). We explored how we can *create the space* for value-centered privacy decisions by applying 4DT, conceptualized privacy decisions in terms of its four dimensions: *self-definition*, the commitments we take on how to be and act in the world, where commitments orientated towards a similar desirable end-state encompass our *values*; *self-realization*, deliberating on our values, forming an intention on how to act (a privacy preference), and acting upon this intention when faced with a data privacy decision; *self-unification*, whether how we have acted is consistent with our values; and *self-constitution*, whether we are willing and able to take on commitments concerning data privacy.

We then explored existing data privacy challenges through this lens. We firstly conceptualized notice fatigue as three different types, varying by the severity of the challenge. First-degree notice fatigue involves failures of *self-unification*, where a user is caught in a *double bind*. In these situations, the user values two values equally and cannot decide in a way that fulfills both, thereby failing to *self-unify*. For example, a user who values both *efficiency* and *control* equally cannot adequately fulfill both values due to the high number of data privacy notices they must make every day. Second-degree notice fatigue is when a user does not act on their values when making a privacy decision, despite it being in accordance with their values to do so. This fails to uphold *self-realization* and *self-unification*.¹⁹⁸ In these instances, the user could be *akratic* – that is, forming an intention (privacy preference) that is not in-line with their values are. However, it is most likely an instance of *weakness of will*, where a user intends to act in a manner consistent with their values but fails to do so. Third-degree notice fatigue is the “apathetic user” phenomenon – where a user has become so overwhelmed by the sheer number of privacy notices that they fail to *self-constitute* (take on commitments pertaining to data privacy). We secondly conceptualized a lack of relevant privacy controls as a failure to *self-realize* and *self-unify*, and in some instances, a *double bind*. Lastly, we understand (inappropriate) nudges as frustrating *self-realization* and *self-unification*. In these cases, nudges are deployed either in a manner that may encourage a privacy decision that is not consistent with a user’s values or nudge users who do not wish to be nudged. Critically, we identify that nudges can be appropriate when willingly entered into as a form of *self-binding* – that is, as a means of helping oneself follow through on their values and commitments.

¹⁹⁸ Except in the case of “lucky akratic,” described in Section 3.2.

Conclusion and Future Directions

Table 7-1: Summary of major thesis contributions, findings, and implications for future research

Major Contributions	Major Findings (Empirical Studies)	Implications for Future Directions
<p>Conceptualize value-centered privacy decisions using the Four-Dimensional Theory of Self-Governance (4DT)</p>	<p>The privacy-value-app relationship is highly individualized and context-dependent, with <i>Use</i>, <i>Control</i>, and <i>Community</i> being quite prevalent. These were also frequently in “double bind-like” tensions with each other.</p> <p>The context-dependent nature of the value and privacy preference relationship, high survey data variability, and participant understandability challenges suggest Theory of Basic Human Values (THBV) method limitations.</p> <p>The results of the interviews suggest that a 4DT approach to data privacy captures the role of value in data privacy decision making because we identify expected phenomena (e.g., the <i>inertia bias</i>, the “apathetic user”). We also identified added insights concerning the role of existing structures (e.g., surveillance capitalism and social media monopolies) in bringing about these phenomena, suggesting structural hindrance to value-centered privacy choices.</p>	<p>Further engagement with underrepresented groups to identify their privacy-relevant values</p> <p>Design and testing of a (smartphone) VcPA with the major values identified here: <i>Control</i>, <i>Community</i>, and <i>Use</i></p> <p>Explore notice presentation with clearer value-privacy preference mapping</p> <p>Test customizable VcPA profiles with initial survey sorting</p> <p>Investigate how to harmonize VcPAs with broader privacy approaches, such as existing and future regulation</p> <p>How to integrate a VcPA into an app store</p>
<p>Design a VcPA</p>	<p>The VcPA helped users download more value-consistent¹⁹⁹ apps.</p> <p>The “suggest alternatives” button and page helped users find value-consistent apps.</p> <p>Many users saw themselves in the provided VcPA profiles.</p> <p>Profiles and resulting selective notices did not seem to capture all relevant values.</p> <p>Apple Privacy Label terminology was difficult for users to understand.</p> <p>It was sometimes difficult for users to compare alternative apps and the alternative apps did not always match the desired app function.</p>	<p>Deploying VcPAs in other privacy decision-making contexts</p> <p>Longitudinal studies fine-tune timing of exploratory notices</p> <p>Designing and testing a “suggest alternatives” feature on the selective notice itself and ensuring better app function matching based on app description rather than app keywords</p>

¹⁹⁹ Where value-consistent apps were defined as apps that had a minimal acceptability coefficient >0.1 with the user’s selected profile. See Section 6.2 for more information on value-consistency and Section 4.2 for more on the minimal acceptability coefficient.

We then used the 4DT-based understanding of value-centered privacy decisions to **establish the usability and effectiveness of the value-centered approach by designing a privacy assistant** to help users make app choices that are more in accordance with their personal values (Table 7-1). To inform the design of a smartphone assistant that creates this space for users, I examined existing PPA technology using a 4DT lens, discussed in Chapter 2. Using insights from this examination, I proposed a value-centered, smartphone privacy assistant (VcPA) to help users make more value-centered decisions at one privacy decision point: smartphone app choices. This VcPA consists of three features: selective notices, exploratory notices, and suggesting alternatives. Selective notices, based on a user's values, aimed to help users *self-bind* and act according to their values (*self-realize* and *self-unify*). Exploratory notices aimed to combat the *inertia bias*, where users may stay in a profile even if it no longer matches their values (and therefore possibly act in a way that is not *self-unifying*). Lastly, the “suggest alternatives” feature aimed to recommend users with alternative apps that are consistent with their values as a means of *self-realizing* and *self-unifying*. I also identified a particular challenge for VcPA design – tuning the timing of exploratory notices to balance the risk of the *inertia bias* against generating notice fatigue.

In **Chapter 4**, I described the design of a mixed-methods study to evaluate and provide greater insight into the value-centered approach to data privacy. The study consisted of three phases. Phase I involved an online survey of values, privacy preferences, and smartphone apps. This survey was a modified version of an established methodology for quantitatively assessing human values, called the Short Schwartz Value Survey (SSVS). This survey is theoretically grounded in the Theory of Basic Human Values (TBHV) (Schwartz, 1992; Schwartz et al., 2012), a theory in cross-cultural psychology that postulates that there are universal human values that motivate our actions. We also used the Apple Privacy Label ontology to ask about a participant's privacy preferences. The survey provided quantitative data scoring values as overall life-guiding principles; scoring values when deciding whether to download a specific app; and binary (yes/no) questions concerning the acceptability of certain privacy preferences. Phase II involved testing a prototype VcPA system informed by Phase I results. To accomplish this, a testing environment – called the Mock App Store (MAS) – was designed to test the VcPA. The MAS is a web interface that replicates certain features of the Apple App Store and includes a “virtual” smartphone to “download” apps. Participants in Phase II were asked to browse the MAS and download apps. The system recorded interactions with the VcPA and the MAS, as well as elicited feedback on VcPA features. To provide further depth in our exploration of the value-privacy relationship, Phase III consisted of follow-up semi-structured interviews with some Phase II participants. These interviews probed participants' values, privacy preferences, and app choices on the MAS as well as in their everyday lives. The three phases were integrated using a process of convergent design (Fetters et al., 2013), centered on two research questions that were posed to guide design and analysis: **RQ1**: *What is the relationship between values and privacy preferences when deciding to download an app, if any?* And **RQ2**: *How useful and effective is a value-centered privacy assistant at helping users make app choices consistent with their values?* The online privacy preference and value survey (Phase I) were primarily aimed at answering **RQ1** and the VcPA user study (Phase II) at **RQ2**, with the interviews (Phase III) containing questions pertaining to both research questions (Table 4-1). Phase I was also completed first and used to inform the design of Phases II and III. Initial survey results informed interview question selection, and the value-privacy profiles for the prototype VcPA were derived from the survey data.

In **Chapter 5**, I then presented and discussed results from this mixed-methods investigation into the first research question, **RQ1**: *how values are involved in privacy decisions* – in particular, app choice. Using both survey and qualitative (interview) data, we found that the privacy-value-app relationship is highly individualized and context-dependent, with *Use*, *Control*, and *Community* being quite prevalent values. These were also frequently in tension with each other, and these tensions tended to resemble *double binds* (“*double bind-like*”). The context-dependent nature of the value and privacy preference relationship, high survey data variability, and participant understandability challenges suggest methods limitations of the Theory of Basic Human Values (THBV). The results of the interviews further suggest that a 4DT approach to data privacy captures the role of value in data privacy decision making because we were able to identify expected phenomena (e.g., the *inertia bias*, the “*apathetic user*”). We also identified added insights concerning the role of existing overarching structures (e.g., surveillance capitalism and social media monopolies) in bringing about these phenomena, suggesting structural hindrance to value-centered privacy choices. The role of these structures in creating value tensions was especially noticeable in tensions that most resembled a *double bind* – for example, tensions involving the value *Community* introduced by existing social media monopolies.

In **Chapter 6**, I described the results from testing a prototype VcPA system with users to answer **RQ2**. The results served as a proof-of-concept that a value-centered privacy assistant, designed using privacy preferences *and* values, could help users when making privacy decisions such as choosing apps. We found that the current VcPA prototype was helpful for a subset of users in this study, with the “suggest alternative apps” feature especially well-received and helpful for users. The VcPA helped users download more value-consistent²⁰⁰ apps. However, some results provided insights into areas of improvement. Users saw themselves in many VcPA profiles, suggesting that profiles could be further personalized or customizable. Profiles and resulting selective notices did not seem to capture all relevant values, as participants reported other (implicitly value-laden) reasons for ignoring selective notices. This suggests that the values on which the profiles and selective notices are based need to be further investigated. Apple Privacy Label terminology was difficult for users to understand, suggesting that selective notices and profiles could be made clearer. It was also sometimes difficult for users to compare alternative apps and the alternative apps did not always match the desired app function. This suggests that the “suggest alternatives” page could improve by integrating it on the selective notice itself or basing alternative recommendations on app descriptions rather than keywords.

Section 7.3 Implications for Future Research

Taken together, these contributions suggest areas of future study for furthering the value-centered approach to privacy – including improving notice presentation and profile design, engaging with diverse voices, harmonizing VcPAs and broader privacy approaches, and expanding VcPAs to other privacy decision-making contexts. These are summarized in Table 7-1.

²⁰⁰ Where value-consistency is defined by matching one’s selected VcPA profile.

7.3.1 Future Research on VcPA Notices

VcPA notice design will be an important area of future work. As we saw in Chapters 5 and 6, the presentation of the Apple privacy label preferences – as well as the values themselves – were challenging for the participants to understand. We suggested that future work could utilize an approach similar to the “Privacy Facts” display described in Kelley et al. (2013) to improve notice clarity and display. Their “Privacy Facts” displays, deployed at the time of app download, include a subset of privacy details²⁰¹ as well as other relevant information for apps, such as reviews.²⁰² Even though Kelley and colleagues were looking at Android permissions, presumably, this should apply to both systems, even if they work somewhat differently and their user interface differs. In addition, as our investigation involves values, we could aim to visually display the collected data alongside the associated values to increase user comprehension of selective and exploratory notices (Figure 7-1). We could also provide greater transparency by indicating whether the values are upheld (+) or violated (-) by sharing that type of data. This could provide greater clarity to users regarding why they are receiving the notices they are receiving, and how the system is understanding the role of values in forming privacy preferences. Designing and testing a “suggest alternatives” feature on the selective notice itself and ensuring better app function matching based on app description rather than app keywords. Longitudinal studies with VcPA users will be a critical area of future research to further tune the timing of exploratory notices, balancing the “exploration” that counters the *inertia bias* with excessive friction that may encourage a value-inconsistent choice.

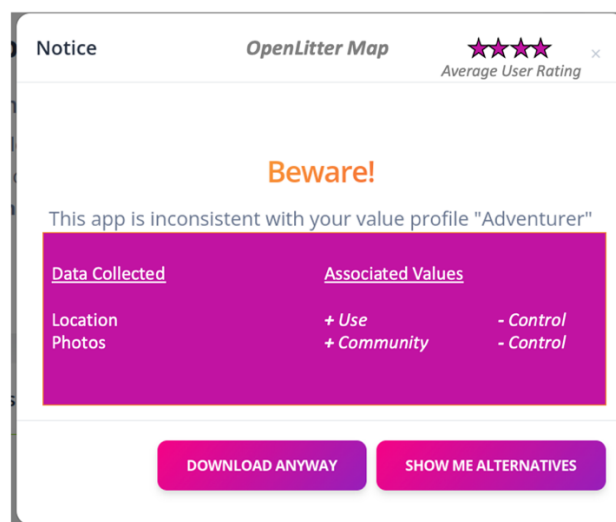


Figure 7-1: Hypothetical example of a selective notice with clearer privacy preference-value mapping and emphasis on values *Use*, *Control*, and *Community*

²⁰¹ These privacy details were selected based on the author’s own previous work as well as Felt et al. (2012), who identified the kinds of data collection and data uses that are most upsetting to users. Intriguingly, Felt and colleagues noted that iOS permissions encompass many of the permissions that ranked highly on their “very upset rate” (VUR) ranking system – such as accessing one’s photos. This (and that Apple has improved their permissions and consent dialogs since 2012) suggests that we should be able to create a more compact list of permissions to include on our VcPA notices.

²⁰² Which were also stated as important to participants in our study – see P04 in Chapter 5 (under the value *Use*) and Mock App Store exit survey comments in Appendix X.

7.3.2 Future Research on VcPA Profiles

Based on user feedback on the VcPA (Chapter 6), we also found that it would be helpful if VcPA profiles could be more tailored to each individual, either by making them more customizable or by using a survey to sort users into profiles. To account for this, future VcPA profiles could begin with a survey to initially “sort” users into a set of primary profiles, allowing them to customize the associated values as they feel fit (Figure 7-1). In addition, many participants in the Mock App Store Study reported seeing themselves in two or more profiles, and further reported that the profiles were not very distinctive. The values should, therefore, not only be on a sliding scale in the survey, but also be presented that way (rather than an all-or-nothing choice) on the profiles themselves. To improve the accuracy of these self-rankings and to try to capture varied interpretations of values, future research could also consider listing each value’s sub-values – the average of which would be the value score.²⁰³




Figure 7-2: Possible future profile design research, with emphasis on profile customizability and the values

Control, Use, and Community

a Engagement with Diverse Voices

It will also be critical for future work to engage with more diverse voices to construct inclusive VcPA profiles. In our study, *all phases* included predominately WEIRD (Western, Educated, Industrial, Rich, and Democratic) participants (Henrich et al., 2010). Participants were predominately well-educated, and well over half in each part of the study had or were in the process of obtaining a Master’s degree (Appendix VI). In each phase, most were of Western (North American and European) nationalities. The interview participants, in particular, were highly educated, sometimes with high levels of knowledge in fields relevant to this study (e.g., philosophy, psychology, HCI, tech, marketing,

²⁰³ For example: basing the ranking for *Community* as the average of sub-value rankings (*accessibility, authenticity, benevolence and universalism, conformity, and connection*).

privacy).²⁰⁴ When studying values in technology, other scholars have claimed that it is important to understand the values of underrepresented and marginalized smartphone user groups to avoid violating their values when designing apps²⁰⁵ (Shams et al., 2023). The same applies to understanding and developing a VcPA system to promote value-centered privacy decisions. An understanding of values based purely on the current WEIRD sample from this study could create a VcPA that minimizes rather than promotes the autonomy of those not represented. The VcPA would fail to consider their values or their value-privacy relationships. The profiles and notices they receive would not be inclusive and relevant to them and could result in nudging them in a direction that is not consistent with their values.²⁰⁶ Engagement with diverse voices will therefore be required to go deeper in the value-privacy-app relationship and to develop a VcPA that is inclusive, beneficial, and ethical.

b Profile Considerations Regarding VcPA Implementation

There are also additional considerations for VcPA profiles when implementing a VcPA in an app store. As indicated in semi-structured interviews (Chapter 5 – particularly *Control* and *Safety*), users are already concerned about the level of information that companies know about them from their data. A VcPA deployed in, say, the Apple App Store, and claiming to understand their “values” may raise concern and resistance.²⁰⁷ Ideally, the logics for a VcPA system – such as their profile and other data – would be stored locally on the user’s phone, interacting with the App Store only to gather information concerning what data the app-in-question will request access to. This, however, would need to be carefully balanced with memory and storage considerations to ensure that the system is not too large.²⁰⁸ One could also consider storing the profile only on the user’s phone, with that alone being sent to the Store where the VcPA logics are – but, due to store browsing associated with an Apple ID, the company would *de facto* get access to the user’s value profile. To respect the choices of all users, including those who do not wish to share their values with an App Store, I would therefore recommend that: 1.) the data obtained from the VcPA, such as a user’s value profile, survey responses, and engagements with notices, *exclusively* be used for the purposes of running the VcPA; 2.) engagement with the VcPA on the app store should be fully voluntary (opt-in); 3.) users must be able to turn the VcPA off; and 4.) all of this information should be clearly communicated when viewing one’s profile. Future translational work into meeting these criteria will be critical for realizing a smartphone VcPA in real-world settings.

²⁰⁴ There are multiple quotes that demonstrate their expertise. For an example, see P08’s reference to Zuboff’s work in Section 5.3.

²⁰⁵ Shams (2023) states that (unintended) value violations by software engineers can be especially harmful to marginalized groups. “Value violations in apps are more destructive if the end-users are vulnerable and marginalized women in conservative societies. [...] For example, recent value violations occurred in 61% (22 out of 36 apps) of menstruation apps, where they shared users’ incredibly personal details with Facebook without the users’ consent [...] This privacy breach is a threat to women’s mental health and might have destructive impacts on their families and social lives” (Shams, 2023, pg. 111648).

²⁰⁶ Thereby violating *self-realization* and *self-unification* according to 4DT. See discussion on inappropriate nudging in Section 3.2

²⁰⁷ “Who are you to tell me what my values are?” (P04, pg. 99)

²⁰⁸ Those who value *utility and function* would likely not like this – see, for example, P23 quote in Section 5.3.

7.3.3 Future Research into Identifying Relevant Values

In order to pursue such future work around VcPA notices and VcPA profiles, we would first need to identify *what* values should be considered. Based on the results presented here (Chapter 5), values are largely context-dependent, rather than general. To account for this complexity, two possible avenues could be used. The first is collaborating with AI to help explore context-dependent values. For example, Liscio et al. (2022) developed a methodology that combines human annotators with AI to understand citizen’s written opinions on various societal issues (e.g., COVID-19). Acknowledging the limitations of Schwartz, one of the aims of their methodology, called *Axies*, was to create *context-specific value lists* for each issue. While using methods such as Liscio et al. (2022) would be interesting in terms of deepening our understanding of value-privacy preference relationships, they note the challenge of defining context boundaries.²⁰⁹ In the case of the VcPA, we could, perhaps, look at creating a VcPA for each category in the app store. However, making so many VcPAs (e.g., VcPA just for finance apps) may not be the most feasible option.²¹⁰

This leads us to our second, and perhaps most promising, area of future research: try re-designing profiles for smartphone-based VcPAs in a manner that supports users in resolving the tension between *Use* and *Control*. As described in Chapter 5, values and privacy preferences seem to be related in a context-dependent manner – but some values, such as *Use* (especially *power and choice*) and *Control* (especially *utility and function*), spanned contexts and participants. Initially building profiles based on these two values could be one possible way forward, as it may help resolve (or, at least, reduce) other value context-specific tensions *by resolving* the frequent tensions between *Use* and *Control*. In particular, the “suggest alternatives” feature of the VcPA – which was widely well-received (but could have better app function match) would allow users to find another app with similar *utility and function* while promoting their *power and choice*. Depending on what kind of app (the context) the user is looking at, other values will be relevant. However, by helping resolve this first tension between *Use and Control* via the “suggest alternatives” feature, the list of alternative apps would hopefully have one that fulfills the user’s greater value set.

There is one important caveat to this idea, however – and this is *Community*. Beside *Control* and *Use*, *Community* (especially *connection*) was also quite prevalent. While *Community* tended to be more relevant in the context of social media and messaging apps (*connection*), other sub-values, such as *conformity*, *authenticity*, and *benevolence and universalism*, were less context constrained. The critical influence of societal and social considerations, such as what app or service one’s friends use, on *connection* also created hindrances to acting according to one’s values. These results have two implications – the first is that profiles based on *Use* and *Control* alone may not be sufficient for helping users act according to their *Community* value, and the second is that an alternative app with similar *connection* value may not exist due to one’s friend network and the market dominance of the app in one’s country. Consider participants like P01, who likes Signal but keeps WhatsApp to connect with friends who are not on Signal (Chapter 5). In these

²⁰⁹ They do, however, note that comparing how distinct the *Axies* results from two different “contexts” are may be one way to determine if the contexts are sufficiently distinct (Liscio et al., 2022, pg. 23). In future work, one could consider using this approach when identifying where different VcPAs are needed.

²¹⁰ We could also consider designing context-specific VcPAs with based on the context-dependent values we identified in the interviews, such as *Safety* for financial apps, and *Pleasure* for entrainment apps. However, like *Axies*, this would be complicated and not the most feasible option.

cases, an alternative app of similar *connection* value is not present, with a strong tension resembling a *double bind* between *Control* and *Community*.

The first implication suggests that *the best area of future research would be to explore building a VcPA based on three values: Control, Use, and Community* (Figures 7-1 and 7-2). This would help capture the *Community* values that are more context-spanning, such as *conformity*. However, *connection* – which is highly related to social media and messaging apps in particular – would not be fully accounted for by this improvement. Given the societal dimension of *connection*, it will likely take some trust-busting and breaking up social media monopolies – or a broader regulatory approach – to fully fulfill this value. Laws mandating data portability and/or cross-platform communication could also be another means for accomplishing this.

7.3.4 Future Research into Harmonizing Regulation and VcPAs

This brings us to a critical insight: the individual value-centered approach to privacy must be complemented by regulatory approaches to further our ideal of respecting user autonomy. As the *connection* example demonstrates, coordination between these two will be critical if we wish to address the tensions (binds) between *Control*, *Use*, and *Community*.²¹¹

As noted in Section 2.3, these regulations are also necessary for upholding collective privacy interests. Privacy has broad, societal implications, and overlaps with other debates on AI transparency, surveillance, and democracies that also must be considered alongside laws that help VcPAs uphold autonomous, value-centered choices (Cohen, 2013; Zuboff, 2019). We saw in the interviews that the systems involved in these debates – such as surveillance capitalism – also influence how much value tensions resemble *double binds* when making a privacy decision. While this could suggest that the individual value-centered approach presented in this thesis is limited like traditional privacy self-management approaches (e.g., notice-and-consent) by nature of its emphasis on the individual, I see it as an opportunity to bridge and harmonize ethical concerns in the individual and collective dimensions. The individual, value-centered privacy approach presented in this thesis could be a complement to regulation, which could set the guardrails of what and how data is collected. In addition, other more global approaches to privacy that better capture *societal norms and values* around privacy, including Contextual Integrity (Nissenbaum, 2004), group privacy (Mittelstadt, 2017), and Privacy-as-Trust (Waldman, 2015), could be employed as a guide for designing regulation, with a value-centered approach deployed at the individual level to capture the realm of *personal values* and preserve user autonomy. Policymakers should also use these approaches and the value-centered approach presented here when crafting new regulations.

Besides informing future regulation, we can consider how VcPAs could be implemented to complement *existing* laws. We can imagine a VcPA designed for IoT devices that complements the upcoming Data Act in the EU (Data Act, 2022). The Data

²¹¹ It is notable that a VcPA *could* be designed to support participants in resolving values tensions that do not have such strong *double bind* characteristics. Recall from Chapter 5 that value tensions tended to more resemble *double binds* when caused (in part or in full) by structural considerations. A VcPA based on these values and their tensions would not need to be harmonized with legislation or other interventions to promote value-centered privacy choices. However, such an approach would likely introduce a high level of complexity given that the more minor values were highly individualized and context-specific (see Chapter 5). In addition, these other values may be accounted for by resolving the *Control* vs. *Use* tension (see the close relationship between *Use* and context-specific values (see Section 7.3.3)). This suggest that deploying VcPAs based on *Control*, *Use*, and *Community* and exploring the means of harmonizing them with regulation are the most promising avenues for future research.

Conclusion and Future Directions

Act was approved in 2023 and, at the time of writing, is set to be enforced in 2025. The Data Act aims to streamline data re-use in order to stimulate innovation and the EU's digital economy. A VcPA designed for IoT devices could serve as an initial layer of data access control for users, encouraging them to pause and reflect upon their relevant values before providing data, while the Data Act would dictate the use of data post-collection. In other words, this would allow for the individual's personal values relevant to the data privacy choice in question to be considered, while the Data Act captures broader considerations in the name of public interest. In addition, the Data Act will work alongside the Digital Market Act (DMA)²¹² to ensure data interoperability between services (Data Act, 2022; Digital Markets Act, 2022), including between different social media and messaging services. This could help with *Community*-related value tensions which, as we saw in the interviews, were largely due to market dominance and social networks exclusive to a single service.

Besides the Data Act and DMA, we have also argued in Doan, Florea, and Carter (2023) that VcPAs could complement the GDPR (General Data Protection Regulation (GDPR), 2016), even in cases where legal processing of data is based on other means than consent.²¹³ In Chapter 2, I presented that the value-centered approach on which the VcPA is normatively rooted in respect for user autonomy. This ethical end, we argued in Doan et al. (2023), is still relevant regardless of the legal basis of consent. We initially explored this idea by considering how VcPAs could be used in collaboration with system-wide technological interventions, such as layered, dynamic consent platforms, to manage complex e-health data flows. We concluded that collaboration between these systems could promote respect for user autonomy (through VcPA-mediated value-centered privacy choices) and meet relevant legal requirements.

Lastly, there are also concerns around the effectiveness of privacy laws that will need to be resolved to impactfully deploy VcPAs. As mentioned in Section 6.3, it remains a challenge to hold app developers accountable for providing correct information (Ali et al., 2023; Jain et al., 2023; Koch et al., 2022; Kollnig et al., 2022; Rodriguez et al., 2023). Because the VcPA is dependent on the accuracy and reliability of the data collection practices disclosed by apps and services, the VcPA will require that laws are enforced to ensure sufficient accountability and transparency. How to overcome these challenges is currently being debated by legal scholars, and the results of these ongoing debates will be critical for eventual VcPA implementation.

7.3.5 Designing and Implementing VcPAs in Other Contexts

In the previous section, we explored designing and deploying VcPAs in a variety of contexts as a complement to regulations. In order to accomplish this, we will require the means of translating the value-centered approach presented here into other critical contexts – such as web “cookie” privacy, data collection conducted by IoT devices, or e-health data sharing. Based on this thesis and the initial VcPA prototype, I have preliminarily collected design questions for constructing VcPAs in other privacy decision-making contexts – tentatively presented here as the *Selective Facilitated Reflection Framework (SFRF)* (Table 7-2). It includes general versions of the features of the VcPA described in Chapter 3 (Table 3-2) with consideration to identifying values such as *Community* that may need both VcPA-based and broader solutions to be promoted. I have also linked each feature back to the relevant 4DT dimensions for completeness – although, as we noted in Chapter 5 when

²¹² The Digital Market Act aims to ensure fair competition in digital spaces by regulating dominant “gatekeeper” companies, such as Alphabet (Google) and Amazon.

²¹³ See footnote 3 for a list of the different bases of processing data under the GDPR.

Conclusion and Future Directions

discussing the interview results, such level of detail may not be necessary to communicate the necessity of each feature for value-centered privacy decision-making. These steps and questions can serve as a tentative starting point for facilitating value-centered choice and designing the VcPAs of the future.

Table 7-2: Tentative framework, the Selective Facilitated Reflection Framework (SFRF), for VcPA design and deployment in other privacy settings

General Feature	Design Questions to Consider	Definition	Relevant 4DT Dimensions	Feature in VcPA for Smartphone App Choice
Selective Friction	<p><i>What method should we use for this friction?</i></p> <p><i>How do we identify the (context-dependent) values for profile creation?</i></p> <p><i>How do we ensure that profiles are sufficiently inclusive of diverse voices?</i></p> <p><i>How do we design profiles and trigger these notices?</i></p> <p><i>How can we make our profiles and notices clear and understandable?</i></p>	Friction deployed in a manner that is personalized to the user's values, acting as a form of <i>self-binding</i> while not overwhelming the user	<p><i>Self-realization</i></p> <p><i>Self-unification</i></p> <p><i>Self-constitution</i></p>	Selective Notices
Exploratory Process	<p><i>How do we balance this exploration process against (unintentional) nudging?</i></p>	A process of mining whether the user's VcPA profile is still relevant to their value set as a means of combating the <i>inertia bias</i>	<p><i>Self-realization</i></p> <p><i>Self-unification</i></p>	Exploratory Notices
Suggest Alternative Action	<p><i>What alternatives do we suggest? How do we select them?</i></p> <p><i>Is harmonization with regulation or a broader approach required to support this alternative action based on related values?</i></p>	To quickly link the user to a relevant alternative that better matches their value set and to relieve value tensions <i>at the point of selective friction</i>	<p><i>Self-realization</i></p> <p><i>Self-unification</i></p>	Suggest Alternative Apps (button and page)

Section 7.4 Concluding Thoughts

To respect user autonomy in data privacy decisions and help them make more meaningful privacy choices, I have proposed here that we design for value-centered privacy decisions. To do this, I have conceptualized value-centered privacy decisions and applied this

Conclusion and Future Directions

understanding to build a prototype privacy assistant that helps users make more value-centered choices.

However, the remedies to the data privacy challenges of our age will require synergy between individual and collective, and regulatory and technological, solutions. The value-centered approach presented here is a means of promoting what makes us unique, empowering us to live a life according to our values in an increasingly data-fueled world. While it is meant to be *empowering*, grounded in our fundamental respect for each other as autonomous agents, it is not *absolute*. Shaping data sharing in a manner that is centered on our own values and the greater norms we wish to base our societies upon means we must come together. This will require dedication on our parts. We must commit to constitute ourselves as autonomous agents and consciously advocate for what we value – not only when making data privacy decisions, but in other spheres as well. We will need to defend our democracies from data misuse and implore our governments to craft effective privacy regulations. To accomplish this, we will need to unite in pursuit of common aims. This will require humility, vulnerability, and a willingness to listen to others and their views on privacy. In brief, it will require citizens, legislators, and technologists working in harmony towards a technological future that we can *all* embrace with enthusiasm.

Bibliography

- 1,675 children removed from parents' custody in benefits scandal. (2022, May 11). *NL Times*.
- Alashoor, T., Keil, M., Liu, L. A., & Smith, J. (2015). How values shape concerns about privacy for self and others. *2015 International Conference on Information Systems: Exploring the Information Frontier*.
- Ali, M. M., Balash, D. G., Kanich, C., & Aviv, A. J. (2023). Honesty is the best policy: On the accuracy of Apple privacy labels compared to apps' privacy policies. *ArXiv*. <http://arxiv.org/abs/2306.17063>
- Allen, A. L. (2013). An ethical duty to protect one's own information privacy? *Faculty Scholarship at Penn Law*, 451.
- Almuhimedi, H., Schaub, F., Sadeh, N., Adjerid, I., Acquisti, A., Gluck, J., Cranor, L., & Agarwal, Y. (2015). Your location has been shared 5,398 times! A field study on mobile app privacy nudging. In K. Inkpen & W. Woo (Eds.), *Proceedings of the 2015 Conference on Human Factors in Computing Systems (CHI)* (pp. 787–796). ACM. <https://doi.org/10.1145/2702123.2702210>
- Apple. (2021, January 27). *Data privacy day at Apple: Improving transparency and empowering users*. Apple Newsroom. <https://www.apple.com/newsroom/2021/01/data-privacy-day-at-apple-improving-transparency-and-empowering-users/?afid=p239%7C10078&cid=aos-us-aff-ir>
- Arvanitis, A., Kalliris, K., & Kaminiotis, K. (2020). Are defaults supportive of autonomy? An examination of nudges under the lens of Self-Determination Theory. *The Social Science Journal*, 1–11. <https://doi.org/https://doi.org/10.1016/j.soscij.2019.08.003>
- Assange, J. (2016). *When Google met Wikileaks*. OR Books.
- Baumard, N., & Sperber, D. (2010). Weird people, yes, but also weird experiments. *Behavioral and Brain Sciences*, 33(2–3), 84–84. <https://doi.org/https://doi.org/10.1017/S0140525X10000038>
- Beauchamp, T. L. (2011). Informed consent: Its history, meaning, and present challenges. *Cambridge Quarterly of Healthcare Ethics*, 20(4), 515–523. <https://doi.org/10.1017/S0963180111000259>
- Ben-Shahar, O., & Schneider, C. E. (2010). The failure of mandated disclosure. In *U of Chicago Law and Economics* (No. 516; Olin Working Paper, Vol. 159, Issue 3). <https://doi.org/10.2139/ssrn.1567284>
- Berlin, I. (1969). Two concepts of liberty. In I. Berlin (Ed.), *Four Essays on Liberty* (pp. 18–172). Oxford University Press.
- Brandt, A. M. (1978). Racism and research: The case of the Tuskegee syphilis study. *The Hastings Center Report*, 8(6), 21. <https://doi.org/10.2307/3561468>
- Braun, V., & Clarke, V. (2022). *Thematic Analysis: A Practical Guide* (A. Maher (Ed.); 1st ed.). SAGE Publications.
- Brignull, H. (n.d.). *Deceptive Design*. Retrieved September 29, 2022, from <https://www.deceptive.design>
- Brumen, B., Zajc, A., & Bošnjak, L. (2023). Permissions vs. privacy policies of apps in Google Play Store and Apple App Store. In M. Tropmann-Frick, H. Aakkola, B. Thalheim, Y. Kiyoki, & N. Yoshida (Eds.), *Information Modelling and Knowledge Bases XXXIV* (pp. 258–275). IOS Press. <https://doi.org/10.3233/FAIA220507>
- Brunton, F., & Nissenbaum, H. (2015). *Obfuscation: A User's Guide for Privacy and Protest*. MIT Press.

Bibliography

- Calo, M. R. (2012). Against notice skepticism in privacy (and elsewhere). *Notre Dame Law Review*, 59(2011), 1027–1072. <https://ssrn.com/abstract=1790144>
- Campbell, J. E., & Carlson, M. (2002). Panopticon.com: Online surveillance and the commodification of privacy. *Journal of Broadcasting and Electronic Media*, 46(4), 586–606. https://doi.org/10.1207/s15506878jobem4604_6
- Caporalet, L. R., & Brewer, M. B. (1991). The quest for human nature: Social and scientific issues in evolutionary psychology. *Journal of Social Issues*, 47(3), 1–9. <https://doi.org/https://doi.org/10.1111/j.1540-4560.1991.tb01819.x>
- Carlson, R. V., Boyd, K. M., & Webb, D. J. (2004). The revision of the Declaration of Helsinki: Past, present and future. *British Journal of Clinical Pharmacology*, 57(6), 695–713. <https://doi.org/10.1111/j.1365-2125.2004.02103.x>
- Cavoukian, A. (2009). Privacy by Design - The 7 foundational principles - Implementation and mapping of fair information practices. *Information and Privacy Commissioner of Ontario, Canada*. <https://doi.org/10.1007/s12394-010-0062-y>
- Childress, J. F. (1990). The place of autonomy in bioethics. *The Hastings Center Report*, 20(1), 12–17. <https://doi.org/10.2307/3562967>
- Chitkara, S., Gothoskar, N., Harish, S., Hong, J., & Agarwal, Y. (2017). Does this app really need my location? Context-Aware privacy management for smartphones. In S. Santini (Ed.), *Proceedings of the Interactive, Mobile, Wearable and Ubiquitous Technologies (UbiComp)* (Vol. 1, Issue 3, pp. 1–22). ACM. <https://doi.org/10.1145/3132029>
- Cohen, J. E. (2013). What privacy is for. *Harvard Law Review*, 126(7), 1904–1933.
- Colnago, J., Feng, Y., Palanivel, T., Pearman, S., Ung, M., Acquisti, A., Cranor, L. F., & Sadeh, N. (2020). Informing the design of a personalized privacy assistant for the Internet of Things. *Proceedings of the 2020 Conference on Human Factors in Computing Systems (CHI)*, 1–13. <https://doi.org/10.1145/3313831.3376389>
- Commissie: ongekend onrecht in toelagenaffaire, beginselen rechtsstaat geschonden. (2020, December 17). NOS. <https://nos.nl/collectie/13855/artikel/2361021-commissie-ongekend-onrecht-in-toelagenaffaire-beginselen-rechtsstaat-geschonden>
- Cox, A. L., Gould, S., Cecchinato, M. E., Iacovides, I., & Renfree, I. (2016). Design frictions for mindful interactions: The case for microboundaries. *Proceedings of the 2016 Conference on Human Factors in Computing Systems (CHI)*, 1389–1397. <https://doi.org/10.1145/2851581.2892410>
- D'Ignazio, C., & Klein, L. F. (2020). *Data Feminism*. MIT Press.
- Das, A., Degeling, M., Smullen, D., & Sadeh, N. (2018). Personalized privacy assistants for the internet of things: Providing users with notice and choice. *IEEE Pervasive Computing*, 17(3), 35–46. <https://doi.org/10.1109/MPRV.2018.03367733>
- Data Act, (2022). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM%3A2022%3A68%3AFIN>
- Davenport, T. H., & Beck, J. C. (2002). *The Attention Economy: Understanding the New Currency of Business* (Revised Ad). Harvard Business Review Press.
- Dayen, D. (2016). Google's remarkably close relationship with Obama's White House. *The Intercept*. <https://theintercept.com/2016/04/22/googles-remarkably-close-relationship-with-the-obama-white-house-in-two-charts/>
- de Wet, J., Wetzelhütter, D., & Bacher, J. (2019). Revisiting the trans-situationality of values in Schwartz's Portrait Values Questionnaire. *Quality and Quantity*, 53(2), 685–711. <https://doi.org/10.1007/s11135-018-0784-8>
- Degeling, M., Utz, C., Lentzsch, C., Hosseini, H., Schaub, F., & Holz, T. (2019). We value your privacy ... now take some cookies: Measuring the GDPR's impact on web privacy. *Symposium on Network and Distributed System Security (NDSS)*.

Bibliography

- <https://doi.org/10.14722/ndss.2019.23378>
- Deibert, R. (2015). Cyberspace under siege. *Journal of Democracy*, 26(3), 64–78.
<https://doi.org/10.1353/jod.2015.0051>
- Derguech, W., Syeda, S. e Z., & D'Aquin, M. (2018). Assessing the readability of policy documents: The case of terms of use of online services. *Proceedings of the 11th International Conference on Theory and Practice of Electronic Governance*, 247–256. <https://doi.org/10.1145/3209415.3209498>
- Digital Markets Act, Pub. L. No. 2022/1925 (2022). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32022R1925&qid=1703171666542>
- Dinno, A. (2015). Nonparametric pairwise multiple comparisons in independent groups using Dunn's test. *Stata Journal*, 15(1), 292–300.
<https://doi.org/10.1177/1536867x1501500117>
- Doan, X., Florea, M., & Carter, S. E. (2023). Legal-Ethical challenges and technological solutions to e-health data consent in the EU. In P. Lukowicz, S. Mayer, J. Koch, J. Shawe-Taylor, & I. Tiddi (Eds.), *HHAI 2023: Augmenting Human Intellect* (pp. 243–253). IOS Press. <https://doi.org/10.3233/FAIA230088>
- Draper, N. A., & Turow, J. (2019). The corporate cultivation of digital resignation. *New Media and Society*, 21(8), 1824–1839. <https://doi.org/10.1177/1461444819833331>
- Dunn, O. J. (1964). Multiple comparisons using rank sums. *Technometrics*, 6(3), 241–252.
- Dupré, J. (2001). *Human Nature and the Limits of Science*. Oxford University Press.
<https://doi.org/https://doi.org/10.1093/0199248060.001.0001>
- Dutch childcare benefit scandal an urgent wake-up call to ban racist algorithms*. (2021). Amnesty International,.
- Dworkin, G. (1988a). Paternalism: Some second thoughts. In *The Theory and Practice of Autonomy* (pp. 121–129). Cambridge University Press.
<https://doi.org/10.1017/cbo9780511625206.009>
- Dworkin, G. (1988b). The nature of autonomy. In *The Theory and Practice of Autonomy* (pp. 1–20). Cambridge University Press. <https://doi.org/10.3402/nstep.v1.28479>
- Estes, A. C. (2011). Some perspective on Obama's bromance with Eric Schmidt. *The Atlantic*. <https://www.theatlantic.com/politics/archive/2011/06/obamas-bromance-google-eric-schmidt-out-hand/352130/>
- General Data Protection Regulation (GDPR), (2016). <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679>
- Felt, A. P., Egelman, S., & Wagner, D. (2012). I've got 99 problems, but vibration ain't one: A survey of smartphone users' concerns. In W. Enck & X. Jiang (Eds.), *Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM)* (pp. 33–43). ACM. <https://doi.org/10.1145/2381934.2381943>
- Ferrer, X., Nuenen, T. Van, Such, J. M., Cote, M., & Criado, N. (2021). Bias and discrimination in AI: A cross-disciplinary perspective. *IEEE Technology and Society Magazine*, 40(2), 72–80. <https://doi.org/10.1109/MTS.2021.3056293>
- Fetters, M. D., Curry, L. A., & Creswell, J. W. (2013). Achieving integration in mixed methods designs - Principles and practices. *Health Services Research*, 48, 2134–2156.
<https://doi.org/10.1111/1475-6773.12117>
- Frankfurt, H. G. (1971). Freedom of the will and the concept of a person. *The Journal of Philosophy*, 68(1), 5–20. <https://doi.org/10.2307/j.ctvvh84xp.15>
- Friedman, B., Kahn, P. H., & Borning, A. (2008). Value sensitive design and information systems. In K. E. Himma & H. T. Tavani (Eds.), *The Handbook of Information and Computer Ethics* (1st ed., pp. 69–103). Wiley.
<https://doi.org/http://dx.doi.org/10.1002/9780470281819.ch4>
- Garrido, G. M., Sedlmeir, J., Uludağ, Ö., Alaoui, I. S., Luckow, A., & Matthes, F. (2022).

Bibliography

- Revealing the landscape of privacy-enhancing technologies in the context of data markets for the IoT: A systematic literature review. *Journal of Network and Computer Applications*, 207. <https://doi.org/10.1016/j.jnca.2022.103465>
- Glaser, B. G., & Strauss, A. L. (1999). *The Discovery of Grounded Theory* (1st ed.). Routledge. <https://doi.org/10.4324/9780203793206>
- Goldhaber, M. H. (1997). The attention economy and the Net. *First Monday*, 2(4). <https://doi.org/10.5210/fm.v2i4.519>
- Grassl, P., Schraffenberger, H., Zuiderveen Borgesius, F., & Buijzen, M. (2021). Dark and bright patterns in cookie consent requests. *Journal of Digital Social Research*, 3(1), 1–38. <https://doi.org/10.31234/osf.io/gqs5h>
- Gray, C. M., Kou, Y., Battles, B., Hoggatt, J., & Toombs, A. L. (2018). The dark (patterns) side of UX design. In A. Dey, E. Cutrell, & M. C. Shruemel (Eds.), *Proceedings of the 2018 Conference on Human Factors in Computing Systems (CHI)* (pp. 1–14). ACM. <https://doi.org/10.1145/3173574.3174108>
- Guttman, L. (1968). A general nonmetric technique for finding the smallest coordinate space for a configuration of points. *Psychometrika*, 33(4), 469–506. <https://doi.org/10.1007/BF02290164>
- Hagendorff, T. (2018). Privacy literacy and its problems. *Journal of Information Ethics*, 27(2), 127–145.
- Harris, T. (2020). *Tristan Harris testifies on Capitol Hill*. Center for Humane Technology. <https://humanetech.com/tristan-harris-testifies-on-capitol-hill-10-min/>
- Hausman, D. M., & Welch, B. (2010). Debate: To nudge or not to nudge. *Journal of Political Philosophy*, 18(1), 123–136. <https://doi.org/10.1111/j.1467-9760.2009.00351.x>
- Hendl, T., Chung, R., & Wild, V. (2020). Pandemic surveillance and racialized subpopulations: Mitigating vulnerabilities in COVID-19 apps. *Journal of Bioethical Inquiry*. <https://doi.org/10.1007/s11673-020-10034-7>
- Henley, J. (2021, January 15). Dutch government resigns over child benefits scandal. *The Guardian*.
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2–3), 61–83. <https://doi.org/10.1017/S0140525X0999152X>
- Heurix, J., Zimmermann, P., Neubauer, T., & Fenz, S. (2015). A taxonomy for privacy enhancing technologies. *Computers and Security*, 53, 1–17. <https://doi.org/10.1016/j.cose.2015.05.002>
- Hoofnagle, C. J., & Urban, J. M. (2014). Alan Westin’s privacy homo economicus. *Wake Forest Law Review*, 14(1), 261–351. <https://doi.org/10.31235/osf.io/ta2z3>
- Jaccard, P. (1901). Étude comparative de la distribution florale dans une portion des Alpes et des Jura. *Bulletin de La Société Vaudoise Des Sciences Naturelles*, 37, 547–579. <http://retro.seals.ch/digbib/view?pid=bsv-002:1901:37::745>
- Jain, A., Rodriguez, D., Alamo, J. M. Del, & Sadeh, N. (2023). ATLAS: Automatically detecting discrepancies between privacy policies and privacy labels. *2023 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, 94–107. <https://doi.org/10.1109/EuroSPW59978.2023.00016>
- Jensen, C., & Potts, C. (2004). Privacy policies as decision-making tools. In M. Tscheligi & E. Dykstra-Erickson (Eds.), *Proceedings of the 2004 Conference on Human Factors in Computing Systems (CHI)* (Vol. 6, Issue 1, pp. 471–478). <https://doi.org/10.1145/985692.985752>
- Jongsma, K. R., Bredenoord, A. L., & Lucivero, F. (2018). Digital medicine: An opportunity to revisit the role of bioethicists. *American Journal of Bioethics*, 18(9),

Bibliography

- 69–70. <https://doi.org/10.1080/15265161.2018.1498952>
- Joseph Turow. (2003). Americans and online privacy: The system is broken. In *Annenberg Public Policy Center Report*.
http://annenbergclassroom.net/Downloads/Information_And_Society/20030701_America_and_Online_Privacy/20030701_online_privacy_report.pdf
- Kahneman, D. (2011). *Thinking, Fast and Slow*. Penguin Books.
- Kant, I. (1959). *Foundations of the metaphysics of morals (Lewis White Beck translation)*. BobbsMerrill Library of Liberal Arts.
- Karadag, E., Kılıçoğlu, G., & Yılmaz-Kılıçoğlu, D. (2018). Cultural validity trouble in measuring value concept: A study on validity of Schwartz Value Survey in Turkish culture. *Cogent Psychology*, 5(1), 1–18.
<https://doi.org/10.1080/23311908.2018.1523517>
- Kelley, P. G., Cranor, L. F., & Sadeh, N. (2013). Privacy as part of the app decision-making process. In S. Bødker, S. Brewster, P. Baudisch, M. Beaudouin-Lafon, & W. E. Mackay (Eds.), *Proceedings of the 2013 Conference on Human Factors in Computing Systems (CHI)* (pp. 3393–3402). ACM.
<https://doi.org/10.1145/2470654.2466466>
- Khan, C. (2021). Is my phone listening to me? We ask the expert. *The Guardian*.
<https://www.theguardian.com/lifeandstyle/2021/oct/29/is-my-phone-listening-to-me-we-ask-the-expert>
- Killmister, S. (2017). *Taking the Measure of Autonomy: A Four-Dimensional Theory of Self-Governance* (1st ed.). Routledge.
<https://doi.org/https://doi.org/10.4324/9781315204932>
- Kim, T. W., & Werbach, K. (2016). More than just a game: Ethical issues in gamification. *Ethics and Information Technology*, 18(2), 157–173. <https://doi.org/10.1007/s10676-016-9401-5>
- Kitchin, R. (2020). *Using digital technologies to tackle the spread of the coronavirus : Panacea or folly?* (No. 44; The Programmable City).
<https://progcity.maynoothuniversity.ie/wp-content/uploads/2020/04/Digital-tech-spread-of-coronavirus-Rob-Kitchin-PC-WP44.pdf>
- Klugman, C. M., Dunn, L. B., Schwartz, J., & Cohen, I. G. (2018). The ethics of smart pills and self-acting devices: Autonomy, truth-telling, and trust at the dawn of digital medicine. *American Journal of Bioethics*, 18(9), 38–47.
<https://doi.org/10.1080/15265161.2018.1498933>
- Koch, S., Wessels, M., Altpeter, B., Olvermann, M., & Johns, M. (2022). Keeping privacy labels honest. *Proceedings on Privacy Enhancing Technologies*, 2022(4), 486–506.
<https://doi.org/10.56553/popets-2022-0119>
- Kollnig, K., Binns, R., Kleek, M. Van, Lyngs, U., Zhao, J., Tinsman, C., & Shadbolt, N. (2021). Before and after GDPR: Tracking in mobile apps. *Internet Policy Review*, 10(4).
- Kollnig, K., Shuba, A., Van Kleek, M., Binns, R., & Shadbolt, N. (2022). Goodbye tracking? Impact of iOS app tracking transparency and privacy labels. *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*.
<https://doi.org/10.1145/3531146.3533116>
- Kruskal, W. H., & Wallis, W. A. (1952). Use of ranks in one-criterion variance analysis. *Journal of the American Statistical Association*, 47(260), 583–621.
<https://doi.org/10.1080/01621459.1952.10483441>
- Leibold, J. (2020). Surveillance in China's Xinjiang region: Ethnic sorting, coercion, and inducement. *Journal of Contemporary China*, 29(121), 46–60.
<https://doi.org/10.1080/10670564.2019.1621529>

Bibliography

- Lewis, M. (2017). *The Undoing Project*. Allen Lane.
- Lindeman, M., & Verkasalo, M. (2010). Measuring values with the Short Schwartz's Value Survey. *Journal of Personality Assessment*, 85(2), 170–178. <https://doi.org/10.1207/s15327752jpa8502>
- Liscio, E., van der Meer, M., Siebert, L. C., Jonker, C. M., & Murukannaiah, P. K. (2022). What values should an agent align with?: An empirical comparison of general and context-specific values. In *Autonomous Agents and Multi-Agent Systems* (Vol. 36, Issue 1). Springer US. <https://doi.org/10.1007/s10458-022-09550-0>
- Liu, B., Andersen, M. S., Schaub, F., Almuhimedi, H., Zhang, S., Sadeh, N., Acquisti, A., & Agarwal, Y. (2016). Follow my recommendations: A personalized privacy assistant for mobile app permissions. In M. E. Zurko, S. Consolvo, & M. Smith (Eds.), *Proceedings of the Twelfth Symposium on Usable Privacy and Security (SOUPS)* (pp. 27–41). USENIX.
- Liu, B., Lin, J., & Sadeh, N. (2014). Reconciling mobile app privacy and usability on smartphones: Could user privacy profiles help? In A. Broder, K. Shim, & T. Suel (Eds.), *Proceedings of the 23rd International Conference on World Wide Web* (pp. 201–211). ACM. <https://doi.org/10.1145/2566486.2568035>
- Luciano, F. (2020). Mind the app—Considerations on the ethical risks of COVID-19 apps. *Philosophy and Technology*, 33(2), 167–172. <https://doi.org/10.1007/s13347-020-00408-5>
- Lucivero, F., & Jongsma, K. R. (2018). A mobile revolution for healthcare? Setting the agenda for bioethics. *Journal of Medical Ethics*, 44(10), 685–689. <https://doi.org/10.1136/medethics-2017-104741>
- Maio, G. R. (2010). Mental representations of social values. In *Advances in Experimental Social Psychology* (Vol. 42, Issue 10, pp. 1–43). [https://doi.org/10.1016/S0065-2601\(10\)42001-8](https://doi.org/10.1016/S0065-2601(10)42001-8)
- Malterud, K., Siersma, V. D., & Guassora, A. D. (2016). Sample size in qualitative interview studies: Guided by information power. *Qualitative Health Research*, 26(13), 1753–1760. <https://doi.org/10.1177/1049732315617444>
- Maseeh, H. I., Nahar, S., Jebarajakirthy, C., Ross, M., Arli, D., Das, M., Rehman, M., & Ashraf, H. A. (2023). Exploring the privacy concerns of smartphone app users: A qualitative approach. *Marketing Intelligence and Planning*, 41(7), 945–969. <https://doi.org/10.1108/MIP-11-2022-0515>
- Mathur, A., Acar, G., Friedman, M. J., Lucherini, E., Mayer, J., Chetty, M., & Narayanan, A. (2019). Dark patterns at scale: Findings from a crawl of 11K shopping websites. In Er. Gilbert & K. Karahalios (Eds.), *Proceedings of the 2019 Conference on Computer Supported Cooperative Work (CSCW)* (Vol. 3, pp. 81:1-81:32). ACM. <https://doi.org/10.1145/3359183>
- McDonald, A. M., & Cranor, L. F. (2008). The Cost of Reading Privacy Policies. *A Journal of Law and Policy for the Information Society*, 4(3), 543–568. <https://kb.osu.edu/handle/1811/72839>
- Meadon, M., & Spurrett, D. (2010). It's not just the subjects – there are too many WEIRD researchers. *Behavioral and Brain Sciences*, 33(2–3), 104–105.
- Mejova, Y., & Kalimeri, K. (2019). Effect of values and technology use on exercise: Implications for personalized behavior change interventions. *Proceedings of the 2019 Conference on User Modeling, Adaptation and Personalization (UMAP)*, 36–45. <https://doi.org/10.1145/3320435.3320451>
- Milligan, G. W., & Cooper, M. C. (1988). A study of standardization of variables in cluster analysis. *Journal of Classification*, 5(2), 181–204. <https://doi.org/10.1007/BF01897163>

Bibliography

- Mitchell-Yellin, B. (2018). Taking the Measure of Autonomy: A Four-Dimensional Theory of Self-Governance. *Notre Dame Philosophical Reviews*.
<https://ndpr.nd.edu/reviews/taking-the-measure-of-autonomy-a-four-dimensional-theory-of-self-governance/>
- Mittelstadt, B. (2017). From individual to group privacy in Big Data analytics. *Philosophy and Technology*, 30(4), 475–494. <https://doi.org/10.1007/s13347-017-0253-7>
- Morley, J., Cowls, J., Taddeo, M., & Floridi, L. (2020). Ethical guidelines for COVID-19 tracing apps. *Nature*, 582, 29–31.
- Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, 79(1), 119–158.
- Nissenbaum, H. (2011). A contextual approach to privacy online. *Daedalus*, 140(4), 32–48.
- Niwattanakul, S., Singthongchai, J., Naenudorn, E., & Wanapu, S. (2013). Using of jaccard coefficient for keywords similarity. *Proceedings of the International MultiConference of Engineers and Computer Scientists*, 1, 380–384.
- Norberg, P. A., Horne, D. R., & Horne, D. A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41(1), 100–126. <https://doi.org/10.1111/j.1745-6606.2006.00070.x>
- Nouwens, M., Liccardi, I., Veale, M., Karger, D., & Kagal, L. (2020). Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. *Proceedings of the 2020 Conference on Human Factors in Computing Systems (CHI)*.
<https://doi.org/10.1145/3313831.3376321>
- Nurwidyanoro, A., Shahin, M., Chaudron, M. R. V., Hussain, W., Shams, R., Perera, H., Oliver, G., & Whittle, J. (2022). Human values in software development artefacts: A case study on issue discussions in three Android applications. *Information and Software Technology*, 141(106731). <https://doi.org/10.1016/j.infsof.2021.106731>
- O’Flaherty, K. (2021, July). All the data WhatsApp and Instagram send to Facebook. *WIRED*.
- O’Neill, C. (2016). *Weapons of Math Destruction* (1st ed.). Broadway Books.
- O’Neill, O. (2002). *Autonomy and Trust in Bioethics*. Cambridge University Press.
<https://doi.org/10.1017/cbo9780511606250>
- O’Neill, O. (2020). Trust and accountability in a digital age. *Philosophy*, 95(1), 3–17.
<https://doi.org/10.1017/S0031819119000457>
- Obar, J. A., & Oeldorf-Hirsch, A. (2018). The biggest lie on the internet: Ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication and Society*, 1–20. <https://doi.org/10.2139/ssrn.2757465>
- Obie, H. O., Hussain, W., Xia, X., Grundy, J., Li, L., Turhan, B., Whittle, J., & Shahin, M. (2021). A first look at human values-violation in app reviews. *International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS)*, 29–38. <https://doi.org/10.1109/icse-seis52602.2021.00012>
- Ombres, D. (2015). NSA domestic surveillance from the Patriot Act to the Freedom Act: The underlying history, constitutional basis, and the efforts at reform. *Seton Hall Legislative Journal*, 39(1), 27–58.
<https://ezproxy.southern.edu/login?url=http%3A%2F%2Fsearch.ebscohost.com%2Flogin.aspx%3Fdirect%3Dtrue%26db%3Da9h%26AN%3D111811359%26site%3Dhost-live%26scope%3Dsite>
- Paci, F., Pizzoli, J., & Zannone, N. (2023). A comprehensive study on third-party user tracking in mobile applications. *Proceedings of the 18th International Conference on Availability, Reliability and Security*. <https://doi.org/10.1145/3600160.3605079>
- Perera, H., Hussain, W., Mougouei, D., Shams, R. A., Nurwidyanoro, A., & Whittle, J.

Bibliography

- (2019). Towards integrating human values into software: Mapping principles and rights of GDPR to values. *Proceedings of the IEEE International Conference on Requirements Engineering*, 404–409. <https://doi.org/10.1109/RE.2019.00053>
- Peters, D., Calvo, R. A., & Ryan, R. M. (2018). Designing for motivation, engagement and wellbeing in digital experience. *Frontiers in Psychology*, 9. <https://doi.org/10.3389/fpsyg.2018.00797>
- Rodriguez, D., Jain, A., Alamo, J. M. del, & Sadeh, N. (2023). Comparing privacy label disclosures of apps published in both the App Store and Google Play Stores. *2023 European Symposium on Security and Privacy Workshops*, 150–157. <https://doi.org/10.1109/EuroSPW59978.2023.00022>
- Roessler, B., & Mokrosinska, D. (2013). Privacy and social interaction. *Philosophy and Social Criticism*, 39(8), 771–791. <https://doi.org/10.1177/0191453713494968>
- Rokeach, M. (1973). *The Nature of Human Values*. Free Press.
- Rosson, M. B., & Carroll, J. (2002). Scenario-Based design. In J. Jacko & A. Sears (Eds.), *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications* (1st ed., pp. 1032–1050). Lawrence Erlbaum Associates. <https://doi.org/10.2307/798660>
- Ryan, K. J., Brady, J. V., Cooke, R. E., Height, D. I., Jonsen, A. R., King, P., Lebacqz, K., Louisell, D. W., Seldin, D. W., Stellar, E., & Turtle, R. H. (1979). *The Belmont Report*. <https://doi.org/10.1021/bi00780a005>
- Ryan, R. M., & Deci, E. L. (2004). Autonomy is no illusion. In J. Greenberg, S. L. Koole, & T. Pyszczynski (Eds.), *Handbook of Experimental Existential Psychology* (1st ed., pp. 455–485). The Guilford Press.
- Sandelowski, M. (1995). Sample size in qualitative research. *Research in Nursing & Health*, 18(2), 179–183. <https://doi.org/https://doi.org/10.1002/nur.4770180211>
- Sandhaus, H. (2023). Promoting bright patterns. *CHI '23: Designing Technology and Policy Simultaneously Workshop*, 1–9. <https://brightpatterns.org/>
- Schaub, F., Balebako, R., Durity, A. L., & Cranor, L. F. (2015). A design space for effective privacy notices. In L. F. Cranor, B. Robert, & S. Consolvo (Eds.), *Proceedings of the 11th Symposium on Usable Privacy and Security (SOUPS 2015)* (pp. 1–17). USEBIX Association. <https://doi.org/10.1017/9781316831960.021>
- Schmidt, A. T., & Engelen, B. (2020). The ethics of nudging: An overview. *Philosophy Compass*, 15(4), 1–13. <https://doi.org/10.1111/phc3.12658>
- Schwartz, S. H. (1992). Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries. *Advances in Experimental Social Psychology*, 25(C), 1–65. [https://doi.org/10.1016/S0065-2601\(08\)60281-6](https://doi.org/10.1016/S0065-2601(08)60281-6)
- Schwartz, S. H. (1994). Are there universal aspects in the structure and contents of human values? *Journal of Social Issues*, 50(4), 19–45. <https://doi.org/10.1111/j.1540-4560.1994.tb01196.x>
- Schwartz, S. H. (2012). An overview of the Schwartz Theory of Basic Values. *Online Readings in Psychology and Culture*, 2(1), 1–20. <https://doi.org/10.9707/2307-0919.1116>
- Schwartz, S. H., & Bilsky, W. (1987). Toward a universal psychological structure of human values. *Journal of Personality and Social Psychology*, 53(3), 550–562. <https://doi.org/10.1037/0022-3514.53.3.550>
- Schwartz, S. H., Cieciuch, J., Vecchione, M., Davidov, E., Fischer, R., Beierlein, C., Ramos, A., Verkasalo, M., Lönnqvist, J. E., Demirutku, K., Dirilen-Gumus, O., & Konty, M. (2012). Refining the theory of basic individual values. *Journal of Personality and Social Psychology*, 103(4), 663–688. <https://doi.org/10.1037/a0029393>

Bibliography

- Shackleton, J. R. (2019). Alexa, Amazon assistant or government informant? *University of Miami Business Law Review*, 27(2), 301–328.
- Shams, R. A., Shahin, M., Oliver, G., Perera, H., Whittle, J., Nurwidiantoro, A., & Hussain, W. (2023). Investigating end-users' values in agriculture mobile applications development: An empirical study on Bangladeshi female farmers. *Journal of Systems and Software*, 200(111648). <https://doi.org/10.1016/j.jss.2023.111648>
- Simon, B. (2005). The return of panopticism: Supervision, subjection and the new surveillance. *Surveillance and Society*, 3(1), 1–20. <https://doi.org/10.24908/ss.v3i1.3317>
- Snowden, E. J. (2019). *Permanent Record* (1st ed.). Metropolitan Books.
- Solove, D. J. (2002). Conceptualizing privacy. *California Law Review*, 90(4), 1087–1155. <https://doi.org/10.2307/3481326>
- Solove, D. J. (2006). A taxonomy of privacy. *University of Pennsylvania Law Review*, 154(3), 467–498.
- Solove, D. J. (2007). “I’ve got nothing to hide” and other misunderstandings of privacy. *San Diego Law Review*, 44, 745–773. <https://ssrn.com/abstract=998565>
- Solove, D. J. (2021). The myth of the privacy paradox. *George Washington Law Review*, 89(1), 1–51. <https://doi.org/10.2139/ssrn.3536265>
- Spiekermann, S., Grossklags, J., & Berendt, B. (2001). E-privacy in 2nd generation E-commerce: Privacy preferences versus actual behavior. In M. P. Wellman & Y. Shoham (Eds.), *Proceedings of the Conference on Electronic Commerce (EC)* (pp. 38–47). ACM.
- Steinberg, S. B. (2017). Sharenting: Children’s privacy in the age of social media. *Emory Law Journal*, 66(4), 839–884. <http://law.emory.edu/>
- Story, P., Cranor, L. F., Smullen, D., Sadeh, N., Acquisti, A., & Schaub, F. (2020). From intent to action: Nudging users towards secure mobile payments. *Proceedings of the 16th Symposium on Usable Privacy and Security (SOUPS)*, 379–416.
- Stoycheff, E., Liu, J., Xu, K., & Wibowo, K. (2019). Privacy and the Panopticon: Online mass surveillance’s deterrence and chilling effects. *New Media and Society*, 21(3), 602–619. <https://doi.org/10.1177/1461444818801317>
- Susser, D. (2019). Notice after notice-and-consent: Why privacy disclosures are valuable even if consent frameworks aren’t. *Journal of Information Policy*, 9, 37–62. <https://doi.org/10.5325/jinfopoli.9.2019.0132>
- Susser, D., Roessler, B., & Nissenbaum, H. (2019). Technology, autonomy, and manipulation. *Internet Policy Review*, 8(2), 1–22. <https://doi.org/10.14763/2019.2.1410>
- Terpstra, A., Schouten, A. P., Rooij, A. de, & Leenes, R. E. (2019). Improving privacy choice through design: How designing for reflection could support privacy self-management. *First Monday*, 24(6). <https://doi.org/https://doi.org/10.5210/fm.v24i7.9358>
- Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving Decisions about Health, Wealth, and Happiness*. Caravan Books.
- ePrivacy Directive, Official Journal of the European Union 11 (2009). <https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2009:337:0011:0036:en:PDF>
- Troullinou, P. (2017). *Exploring the Subjective Experience of Everyday Surveillance: The Case of Smartphone Devices as Means of Facilitating "Seductive" Surveillance*. The Open University.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: heuristics and biases. *Science*, 185(4157), 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>
- Usman, M. (2022). Breaking up Big Tech: Lessons from AT&T. *University of*

Bibliography

- Pennsylvania Law Review*, 170(2), 523–548.
- Utz, C., Degeling, M., Fahl, S., Schaub, F., & Holz, T. (2019). (Un)informed consent: Studying GDPR consent notices in the field. In D. M. Freeman, A. Mitrokotsa, & A. Sinha (Eds.), *Proceedings of the 2019 Conference on Computer and Communications Security* (pp. 973–990). ACM. <https://doi.org/10.1145/3319535.3354212>
- van Ooijen, I., Segijn, C. M., & Oprea, S. J. (2022). Privacy cynicism and its role in privacy decision-making. In *Communication Research*. Sage. <https://doi.org/10.1177/00936502211060984>
- Velazco, C. (2022). What your Android phone’s new “data safety” labels mean. *Washington Post*.
- Wakefield, J. (2022). People devote third of waking time to mobile apps. *BBC*.
- Waldman, A. E. (2015). Privacy as trust: Sharing personal information in a networked world. *University of Miami Law Review*, 69(3), 559–630. <https://doi.org/10.2139/ssrn.2309632>
- Waldman, A. E. (2018). Privacy, notice, and design. *Stanford Technology Law Review Tech Law Review*, 21(1), 74–127. <https://doi.org/10.2139/ssrn.2780305>
- Warberg, L., Acquisti, A., & Sicker, D. (2019). Can privacy nudges be tailored to individuals’ decision making and personality traits? In J. Domingo-Ferrer (Ed.), *Workshop on Privacy and Electronic Society (WPES)* (pp. 175–197). ACM. <https://doi.org/10.1145/3338498.3358656>
- WhatsApp penetration rate among global messaging app users as of April 2022, by country*. (2023, September 25). Statista.
- World Medical Association. (2013). Declaration of Helsinki, ethical principles for scientific requirements and research protocols. In *Bulletin of the World Health Organization* (Vol. 79, Issue 4). <https://www.wma.net/policies-post/wma-declaration-of-helsinki-ethical-principles-for-medical-research-involving-human-subjects/>
- Yeung, K. (2017). ‘Hyper-nudge’: Big Data as a mode of regulation by design. *Information Communication and Society*, 20(1), 118–136. <https://doi.org/10.1080/1369118X.2016.1186713>
- Zhang, S., Feng, Y., Yao, Y., Cranor, L. F., & Sadeh, N. (2022). How usable are iOS app privacy labels? *Proceedings on Privacy Enhancing Technologies*, 4, 204–228. <https://doi.org/10.56553/popets-2022-0106>
- Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (1st ed.). Profile Books.
- Zuiderveen Borgesius, F. J. (2015). *Improving Privacy Protection in the Area of Behavioural Targeting*. Kluwer Law International. <https://papers.ssrn.com/abstract=2654213>

Appendices

Appendix I: Online Value and Privacy Preference Survey

Questions 1-13 are the consent form and are excluded here. There were two versions of the survey. Questions 14-30 were on both versions of the survey, while questions 30-44 were either pertaining to Lose It! (version 1) or OpenLitterMap (version 2).

Privacy, Values, and Smartphones



Thank you for your interest in privacy, values, smartphones!

The purpose of this study is to better understand how our personal values are involved in our data privacy choices.

In this survey, you will be asked a series of questions about your personal values and privacy preferences when selecting smartphone applications.

It should only take a maximum of ten minutes to complete.

Please note that all data collected from this survey will be aggregated and held in a confidential manner.

Required

Part 1 - Your Demographic Information

14. Please select your age *

- 18-24
- 25-64
- 65 and over
- Prefer not to say

15. Please select your gender *

- Female
- Male
- Other
- Prefer not to say

16. Please select the region that best fits your nationality (listed on your passport) *

- North America (example: United States)
- South America (example: Brazil)
- Asia (example: China)
- Europe (example: the Netherlands)
- Other (example: Australia)
- Prefer not to say

17. Please select your highest level of education. If you are currently enrolled in a degree program, please select that degree. *

- Secondary Degree, High School, or Equivalent
- Bachelor's Degree or Equivalent
- Master's Degree or Equivalent
- Doctorate (PhD, MD) or Equivalent
- Prefer not to say

18. Do you consider yourself a native or fluent English speaker? *

- Yes - a native English speaker
- Yes - a fluent English speaker
- No - not a native or fluent English speaker

19. Do you own or have you previously owned a smartphone? *

- Yes - I currently own a smartphone
- Yes - I do not own a smartphone now, but I have previously owned a smartphone
- No - I do not own or have previously owned a smartphone

Appendix

Part 2 – Your Personal Values

This section focuses on your personal values.

20. Please, rate the overall importance of **social power, authority, and wealth** as life-guiding principles for you. These principles are part of the personal value: **Power**.

Please use the following 9-point scale ranging from 1 (opposed to your principles) to 9 (of supreme importance to you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Opposed Of supreme importance

21. Please, rate the overall importance of **being successful, ambitious, and influential** as life-guiding principles for you. These principles are part of the personal value: **Achievement**.

Please use the following 9-point scale ranging from 1 (opposed to your principles) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Opposed Of supreme importance

22. Please, rate the overall importance of **gratifying your desires, enjoying life, and self-indulgence** as life-guiding principles for you. These principles are part of the personal value: **Hedonism**.

Please use the following 9-point scale ranging from 1 (opposed to your principles) to 9 (of supreme importance to me).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Opposed Of supreme importance

23. Please, rate the overall importance of **living a varied, challenging and exciting life** as life-guiding principles for you. These principles are part of the personal value: **Stimulation**.

Please use the following 9-point scale ranging from 1 (opposed to your principles) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Opposed Of supreme importance

24. Please, rate the overall importance of **creativity, freedom, curiosity, independence, and choosing your own goals** as life-guiding principles for you. These principles are part of the personal value: **Self-Direction**.

Please use the following 9-point scale ranging from 1 (opposed to your principles) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Opposed Of supreme importance

26. Please, rate the overall importance of **being helpful, honest, forgiving, loyal, and responsible** as life-guiding principles for you. These principles are part of the personal value: **Benevolence**.

Please use the following 9-point scale ranging from 1 (opposed to your principles) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Opposed Of supreme importance

25. Please, rate the overall importance of **being broad-minded and wise; pursuing social justice, a world at peace, or equality; and protecting the environment** as life-guiding principles for you. These principles are part of the personal value: **Universalism**.

Please use the following 9-point scale ranging from 1 (opposed to your principles) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Opposed Of supreme importance

27. Please, rate the overall importance of being **humble, modest, and respecting tradition** as life-guiding principles for you. These principles are part of the personal value: **Tradition**.

Please use the following 9-point scale ranging from 1 (opposed to your principles) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Opposed Of supreme importance

Appendix

28. Please, rate the overall importance of **being obedient, honoring your parents and others, self-discipline, and being polite** as life-guiding principles for you. These principles are part of the personal value: **Conformity**.

Please use the following 9-point scale ranging from 1 (opposed to your principles) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Opposed

Of supreme importance

29. Please, rate the overall importance of **national security, family security, personal security, and social order** as life-guiding principles for you. These principles are part of the personal value: **Security**.

Please use the following 9-point scale ranging from 1 (opposed to your principles) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Opposed

Of supreme importance

Part 2 – Personal Values (continued)

30. Do you think your personal values play a role when selecting a smartphone application? *

- Yes
- No
- Sometimes
- I don't know

Version 1 (Lose It!) presented here. Version 2 would be the same, only replacing “Lose It! with “OpenLitterMap.”

Part 3 - Your Values and Smartphone Applications

This section explores your personal values when deciding whether or not to download a smartphone application.

Please consider the follow app:

Lose It! Is a health and wellness app that helps you track your eating habits and exercise.

31. Do you think you would download Lose It! ? *

- Yes
- No
- Maybe

32. Please, rate the overall importance of **power, authority, and wealth** on your decision whether or not to download Lose It!. These principles are part of the personal value: **Power**.

Please use the following 9-point scale ranging from 1 (not applicable) to 9 (of supreme importance to you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Not applicable

Of supreme importance

33. Please, rate the overall importance of **being successful, ambitious, and influential on people and events** on your decision whether or not to download Lose It!. These principles are part of the personal value: **Achievement**.

Please use the following 9-point scale ranging from 1 (not applicable) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Not applicable

Of supreme importance

34. Please, rate the overall importance of **gratifying your desires, enjoying life, and self-indulgence** on your decision whether or not to download Lose It!. These principles are part of the personal value: **Hedonism**.

Please use the following 9-point scale ranging from 1 (not applicable) to 9 (of supreme importance to me).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Not applicable

Of supreme importance

Appendix

35. Please, rate the overall importance of **living a varied, challenging exciting life** on your decision whether or not to download Lose It!. These principles are part of the personal value: **Stimulation**.

Please use the following 9-point scale ranging from 1 (not applicable) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Not applicable

Of supreme importance

36. Please, rate the overall importance of **creativity, freedom, curiosity, independence, and choosing your own goals** on your decision whether or not to download Lose It!. These principles are part of the personal value: **Self-Direction**.

Please use the following 9-point scale ranging from 1 (opposed to your principles) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Not applicable

Of supreme importance

37. Please, rate the overall importance of **being broad-minded and wise; pursuing social justice, a world at peace, or equality; and protecting the environment** on your decision whether or not to download Lose It!. These principles are part of the personal value: **Universalism**.

Please use the following 9-point scale ranging from 1 (not applicable) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Not applicable

Of supreme importance to you

38. Please, rate the overall importance of **being helpful, honest, forgiving, loyal, and responsible** on your decision whether or not to download Lose It!. These principles are part of the personal value: **Benevolence**.

Please use the following 9-point scale ranging from 1 (not applicable) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Not applicable

Of supreme importance

39. Please, rate the overall importance of being **humble, modest, and respecting tradition** on your decision whether or not to download Lose It!. These principles are part of the personal value: **Tradition**.

Please use the following 9-point scale ranging from 1 (not applicable) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Not applicable

Of supreme importance

40. Please, rate the overall importance of **being obedient, honoring your parents and others, self-discipline, and being polite** on your decision whether or not to download Lose It!. These principles are part of the personal value: **Conformity**.

Please use the following 9-point scale ranging from 1 (not applicable) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Not applicable

Of supreme importance

41. Please, rate the overall importance of **national security, family security, personal security, and social order** on your decision whether or not to download Lose It!. These principles are part of the personal value: **Security**.

Please use the following 9-point scale ranging from 1 (not applicable) to 9 (of supreme importance you).

*

1	2	3	4	5	6	7	8	9
---	---	---	---	---	---	---	---	---

Not applicable

Of supreme importance

Appendix

Part 4 - Your Privacy Preferences and Smartphone Applications

This section explores your privacy preferences for your smartphone apps.

For the following questions, imagine that you've downloaded the Lose It! app and it would like to collect data from your device.

Reminder: Lose It! Is a health and wellness app that helps you track your eating habits and exercise.

42. If the data collected **WILL NOT** be linked to you ("unlinked" data), which data would you feel comfortable sharing with Lose It? You may select multiple answers.

In this case, *unlinked* data is data collected by the Lose It! app and then stripped of your name and/or your phone's unique identifier. While these steps are taken to protect your identity, you could still, in theory, be identified if the data is combined with other data sets. *

- Your **health and fitness information** (such as your health data and exercise data)
- Your **financial information** (such as your credit card number, bank account number, form of payment, credit score, salary, income, and debts)
- Your **location**
- Your **sensitive information** (such as your race, ethnicity, religion, political affiliation, or sexual orientation)
- Your **contacts** (such as your address book)
- The **content of your phone** (such as emails, text message, photos, and audio data)
- Your **browsing history** (such as when you are browsing a website, outside of the app)
- Your **search history** in the app
- Your **purchase history** in the app
- Your **usage data** (such as clicks, scrolls, taps, or advertisement views)
- General **diagnostic data** (such as crash logs, launch time, and energy use)
- None

43. If the data collected **WILL** be linked to you, which data would you feel comfortable sharing with Lose It? You may select multiple answers. *

- Your **contact information** (such as your name, email address, phone number, or physical address)
- Your **health and fitness information** (such as your health data and exercise data)
- Your **financial information** (such as your credit card number, bank account number, form of payment, credit score, salary, income, and debts)
- Your **location**
- Your **sensitive information** (such as your race, ethnicity, religion, political affiliation, or sexual orientation)
- Your **contacts** (such as your address book)
- The **content of your phone** (such as emails, text message, photos, and audio data)
- Your **browsing history** (such as when you are browsing a website, outside of the app)
- Your **search history** in the app
- Your **purchase history** in the app
- Other identifiers** (such as your account username or the identifier for your phone)
- Your **usage data** (such as clicks, scrolls, taps, or advertisement views)
- General **diagnostic data** (such as crash logs, launch time, and energy use)
- None


Appendix

44. If the data collected from you **was used for tracking**, which data would you allow Lose It! to collect?

In this case, *tracking* is linking data collected from you in the Lose It! app with data from other apps or websites. This is often done to target you with specific advertisements. *

- Your **contact information** (such as your name, email address, phone number, or physical address)
- Your **health and fitness information** (such as your health data and exercise data)
- Your **financial information** (such as your credit card number, bank account number, form of payment, credit score, salary, income, and debts)
- Your **location**
- Your **sensitive information** (such as your race, ethnicity, religion, political affiliation, or sexual orientation)
- Your **contacts** (such as your address book)
- The **content of your phone** (such as emails, text message, photos, and audio data)
- Your **browsing history** (such as when you are browsing a website, outside of the app)
- Your **search history** in the app
- Other **identifiers** (such as your account username or the identifier for your phone)
- Your **purchase history** in the app
- Your **usage data** (such as clicks, scrolls, taps, or advertisement views)
- General **diagnostic data** (such as crash logs, launch time, and energy use)
- None

This content is neither created nor endorsed by Microsoft. The data you submit will be sent to the form owner.



Appendix

Appendix II: Apps in Mock App Store

Grouped by similar app family and includes acceptability coefficients.

Title	Family	Coefficient Adventurer	Coefficient Goal Setter	Coefficient Helpful Neighbor
STRETCHIT: Stretching Mobility	1	0.20165289256198300	0.2336283185840710	0.15384615384615400
BetterMe: Period Tracker	1	0.19716646989374300	0.21238938053097300	0.16117216117216100
Onyx: Home Workout	1	0.24439197166469900	0.2781289506953220	0.23076923076923100
Relish: Relationship & Couples	1	0.25	0.24557522123893800	0.19230769230769200
Coral: Relationship self-care	1	0.20798898071625300	0.23303834808259600	0.1752136752136750
Impulse - Brain Training	1	0.2165289256198350	0.22300884955752200	0.158974358974359
Lasting: Marriage & Couples	1	0.21028466483011900	0.23795476892822000	0.19373219373219400
Avrora - Sleep Booster	1	0.13636363636363600	0.17256637168141600	0.1282051282051280
Paired: Couples & Relationship	1	0.20425029515938600	0.21491782553729500	0.15018315018315000
Alive by Whitney Simmons	2	0.23037190082644600	0.26991150442477900	0.22435897435897400
Paceline: Rewards for Exercise	2	0.13459268004722600	0.1883691529709230	0.1575091575091580
1st Phorm	2	0.16115702479338800	0.24778761061946900	0.1923076923076920
Achievement - Reward Health	3	0.19952774498229000	0.2262958280657400	0.1758241758241760
Nutrition Coach - Food tracker	3	0.4325068870523420	0.471976401179941	0.376068376068376
Heart Rate Monitor - Pulse HR	3	0.045454545454545500	0.048672566371681400	0.038461538461538500
Organic Fit: Home Weight Loss	3	0.15220385674931100	0.193952802359882	0.14957264957265000
Lunar - Period Tracker	3	0.3181818181818180	0.31194690265486700	0.23076923076923100
MindFull: Weight Loss Hypnosis	3	1.0	1.0	1.0
Drink Water Reminder, Tracker	3	1.0	1.0	1.0
K Health Telehealth	3	0.11983471074380200	0.1592920353982300	0.12051282051282100
Sleep Monitor: Sleep Recorder	3	0.1971664698937430	0.19974715549936800	0.1391941391941390
Mintal Tracker: Sleep Recorder	3	0.20000000000000000	0.2	0.13846153846153800
Ovulation Calculator Fertile Tracker & Calendar OC	4	1.0	1.0	1.0
Glow Period, Fertility Tracker	4	0.15805785123966900	0.19634955752212400	0.15224358974359000
Ovia Fertility & Cycle Tracker	4	0.15922865013774100	0.20058997050147500	0.158974358974359
ShutEye: Sleep Tracker	4	0.25757575757575800	0.2846607669616520	0.21794871794871800
Period Tracker My Cycle	4	0.29476584022038600	0.28613569321533900	0.20512820512820500
Life - Period Tracker Calendar	4	0.3181818181818180	0.31194690265486700	0.23076923076923100
Clover Period Tracker Calendar	4	0.1277813095994910	0.16269571136827800	0.1301775147928990
Premom Ovulation Tracker	4	0.1487603305785120	0.19690265486725700	0.16025641025641000
Flo Period & Ovulation Tracker	4	0.1351829988193630	0.17951959544879900	0.13736263736263700
Menstrual Cycle Tracker	4	1.0	1.0	1.0
Fertility & Period Tracker	4	1.0	1.0	1.0
Pregnancy Test Checker	4	0.11735537190082600	0.16460176991150400	0.09743589743589740
Eat This Much - Meal Planner	5	0.2628099173553720	0.2920353982300890	0.24102564102564100
MyPlate Calorie Counter	5	0.3360881542699720	0.359882005899705	0.32478632478632500

Appendix

MyFitnessPal	5	0.19375573921028500	0.22812192723697100	0.18803418803418800
ControlMyWeight	5	0.12396694214876000	0.17699115044247800	0.13846153846153800
Calory: Simple Calorie Counter	5	0.2165289256198350	0.22300884955752200	0.158974358974359
Gymshark Training: Fitness App	5	0.11129476584022000	0.16342182890855500	0.12136752136752100
Calorie Counter +	5	0.15151515151515200	0.19764011799410000	0.1623931623931620
Lose It!	5	0.14639905548996500	0.19974715549936800	0.16483516483516500
Calorie Counter - MyNetDiary	5	0.15151515151515200	0.19764011799410000	0.1623931623931620
Macros - Calorie Counter	5	0.21605667060212500	0.24778761061946900	0.18681318681318700
Gratitude Journal Affirmations	6	0.16765053128689500	0.17572692793931700	0.11721611721611700
Manifest - Affirmations	6	0.14951164537941400	0.16975060337892200	0.12354312354312400
Affirm It	6	0.2066115702479340	0.2168141592920350	0.17307692307692300
Jour: Daily Self-Care Journal	6	0.2190082644628100	0.24778761061946900	0.20512820512820500
ThinkUp - Daily Affirmations	6	0.2243211334120430	0.24399494310998700	0.18681318681318700
Shine: Calm Anxiety & Stress	6	0.28264462809917400	0.3256637168141590	0.26153846153846200
Bloom: CBT Therapy & Self-Care	6	0.1806375442739080	0.22123893805309700	0.1575091575091580
#Mindful - Motivation Quotes	6	0.29476584022038600	0.28613569321533900	0.20512820512820500
Quo: Daily Motivation Quotes	6	1.0	1.0	1.0
Sanity & Self: Stress Relief	6	1.0	1.0	1.0
Motivate: Daily Motivation	6	0.1928374655647380	0.26843657817109100	0.21367521367521400
Motivation - Daily quotes	6	0.12231404958677700	0.16460176991150400	0.12051282051282100
I am - Daily Affirmations	6	0.120961682945154	0.16411906677393400	0.11888111888111900
Quotes: Daily Inspiration	6	0.20425029515938600	0.21491782553729500	0.15018315018315000
Mantra - Daily Affirmations	6	0.1460055096418730	0.19616519174041300	0.11111111111111110
Depression Test	6	0.30303030303030303	0.3008849557522120	0.20512820512820500
Happify: for Stress & Worry	7	0.19338842975206600	0.23185840707964600	0.18205128205128200
Meditation & Sleep Mindfulness	7	0.23494687131050800	0.2667509481668770	0.2161172161172160
Breath: Meditation & Sleep	7	0.25	0.2676991150442480	0.1987179487179490
MindDoc: Your Companion	7	0.18016528925619800	0.21061946902654900	0.1794871794871800
The Mindfulness App	7	0.13957759412304900	0.19174041297935100	0.150997150997151
Meditation Studio	7	0.131198347107438	0.17035398230088500	0.1346153846153850
7Mind Meditation & Sleep	7	0.22004132231405000	0.23893805309734500	0.17628205128205100
Relaxing Sounds, Sleep Easy	7	1.0	1.0	1.0
Slumber: Fall Asleep, Insomnia	7	0.1983471074380170	0.2433628318584070	0.20512820512820500
Soothing Sleep Sounds	7	0.2066115702479340	0.25663716814159300	0.21367521367521400
Calm: Sleep & Meditation	7	0.15537190082644600	0.2097345132743360	0.15897435897435900
Headspace: Mindful Meditation	7	0.15794306703397600	0.21140609636184900	0.16524216524216500
Smiling Mind	7	0.15702479338843000	0.21902654867256600	0.15384615384615400
Sleep Sounds by Sleep Pillow	7	0.4421487603305790	0.42920353982300900	0.3076923076923080
Simple Habit Sleep, Meditation	8	0.19008264462809900	0.23205506391347100	0.18803418803418800
Alo Moves	8	0.6033057851239670	0.504424778761062	0.46153846153846200

Appendix

Lotus Yoga and Workout	8	0.22644628099173600	0.2584070796460180	0.2153846153846150
Simply Yoga - Home Instructor	8	0.22561983471074400	0.2654867256637170	0.22564102564102600
5 Minute Yoga Workouts	8	1.0	1.0	1.0
Oak - Meditation & Breathing	8	1.0	1.0	1.0
Yoga with Gotta Joga	8	0.09090909090909090	0.18584070796460200	0.10256410256410300
Yogaia: Inspiring workouts	8	0.15702479338843000	0.2157079646017700	0.16987179487179500
Yoga Down Dog	8	0.14214876033057900	0.19646017699115000	0.15384615384615400
Glo Yoga and Meditation App	8	0.16942148760330600	0.22713864306784700	0.16666666666666700
Yoga - Poses & Classes at Home	8	0.1322314049586780	0.17625368731563400	0.1474358974358970
Yoga for Beginners Mind+Body	8	0.174931129476584	0.22271386430678500	0.18376068376068400
MyLife Meditation: Mindfulness	8	0.1334120425029520	0.19848293299620700	0.15750915750915800
MindShift CBT - Anxiety Relief	8	1.0	1.0	1.0
O My Soul Christian Meditation	9	0.20000000000000000	0.2	0.13846153846153800
Soultime Christian Meditation	9	0.20543093270366000	0.24020227560050600	0.1904761904761910
Yoga Studio: At-home classes	9	0.18417945690673000	0.1972187104930470	0.16849816849816800
Abide: Pray & Relax Meditation	9	0.17768595041322300	0.19469026548672600	0.13675213675213700
Pray.com: Prayer, Sleep, Bible	9	0.1168831168831170	0.163716814159292	0.12271062271062300
SlimFast Together	9	0.15702479338843000	0.2153392330383480	0.1752136752136750
Sleep Sounds & Relax: MindZone	9	0.26859504132231400	0.2654867256637170	0.19658119658119700
Prayer Guide - Bible Devotions	9	0.19421487603305800	0.25663716814159300	0.1282051282051280
SilverSneakers GO	9	0.16942148760330600	0.2278761061946900	0.16666666666666700

Appendix

Appendix III: Entrance Survey

Entrance Survey - Mock App Store



Thank you for your interest in privacy, values, smartphones!

The purpose of this study is to better understand how our personal values are involved in our data privacy choices.

For the **first stage of this study**, please complete a short entrance survey. This survey will ask you about your demographic details and preliminarily assess your privacy preferences. Please note that all data collected from this survey will be held in a confidential manner.

This survey should only take a minute to complete. The entire study will take about 15 minutes and must be completed all at once. **Please make sure to click the link at the end of the survey to continue on to the rest of the Study!**

After you complete this survey, you will be directed to the second stage of the study.

required

Part 1 - Your Demographic Information

1. Please select your age *

- 18 - 24
- 24 - 64
- 64 and over
- Prefer not to say

2. Please select your gender *

- Female
- Male
- Other
- Prefer not to say

3. Please select the region that best fits your nationality (listed on your passport) *

- North America (example: United States)
- South America (example: Brazil)
- Asia (example: China)
- Europe (example: the Netherlands)
- Other (examples: Nigeria or Australia)
- Prefer not to say

4. Please select your highest level of education. If you are currently enrolled in a degree program, please select that degree. *

- Secondary Degree, High School, or Equivalent
- Bachelor's Degree or Equivalent
- Master's Degree or Equivalent
- Doctorate (PhD, MD) or Equivalent
- Prefer not to say

5. Do you consider yourself a native or fluent English speaker? *

- Yes - a native English speaker
- Yes - a fluent English speaker
- No - not a native or fluent English speaker

6. Do you own or have you previously owned a smartphone? *

- Yes - I currently own a smartphone
- Yes - I do not own a smartphone now, but I have previously owned a smartphone
- No - I do not own or have previously owned a smartphone

Appendix

Part 2 - Your Views on Data Privacy

This section asks you three general questions about how you view data privacy.

7. How concerned are you about your data privacy *in general*? *

1	2	3	4	5
---	---	---	---	---

Not at all concerned Very concerned

8. How concerned are you about your data privacy *on your smartphone*? *

1	2	3	4	5
---	---	---	---	---

Not at all concerned Very concerned

9. How useful do you find the privacy notifications and controls on your current smartphone? *


1	2	3	4	5
---	---	---	---	---

Not at all useful Very useful

Thank you for your interest in this study.

Based on your responses, you do not qualify for this study. Thank you for your time and interest. **Please close this window to exit the form. To not press submit.**

This content is neither created nor endorsed by Microsoft. The data you submit will be sent to the form owner.



Appendix

Appendix IV: Exit Survey

Exit Survey - Mock App Store

Thank you for completing the "Mock App Store" exercise!

Please complete this short exit survey on your experience.

It should only take about 2 minutes to complete.

Please note that all data collected from this survey will be aggregated and held in a confidential manner.

* Required

Your Feedback on the Mock App Store Experience

Please provide feedback on your experience.

1. In general, how helpful did you find **the notifications** that you received? *

1 2 3 4 5

Not at all helpful Very helpful

2. How helpful were the **timing** of pop-up notifications? *

1 2 3 4 5

Not at all helpful Very helpful

3. How did you find the **frequency** of the notifications? *

1 2 3 4 5

Too little Too many

4. How helpful did you find the **"show me alternative applications"** feature? *

1 2 3 4 5

Not at all helpful Very helpful

5. *Optional:* What other **features or information** would help you select a smartphone application?

Please specify briefly:

6. During the Mock App Store Exercise, how confident are you that you made app choices consistent with your values? *

1 2 3 4 5

Not at all confident Very confident

9. How concerned are you about data privacy *on your smartphone*? *

1 2 3 4 5

Not at all concerned Very concerned

7. Please rank the following statements about value profiles. *

	False	Somewhat False	Neutral	Somewhat True	True
I found the profiles clear.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I found a profile that matched my values well.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The values I care about were listed in a profile.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

8. How concerned are you about your data privacy *in general*? *

1 2 3 4 5

Not at all concerned Very concerned

This content is neither created nor endorsed by Microsoft. The data you submit will be sent to the form owner.



Appendix

Appendix V: Pre-Interview Survey

Pre-Interview Survey and Information

Thank you for your interest in privacy, values, smartphones!

The purpose of this study is to better understand how our personal values are involved in our data privacy choices.

Before your interview, please complete this **short 6-question survey**. This survey will ask you about your demographic details.

Please note that all data collected from this survey will be reported in aggregate and held in a confidential manner.

It should only take a minute to complete.

* required

Part 1 - Your Demographic Information

1. Please select your age *

- 18 - 24
- 25 - 64
- 65 and over
- Prefer not to say

2. Please select your gender *

- Female
- Male
- Other
- Prefer not to say

3. Please select the region that best fits your nationality (listed on your passport) *

- North America (example: United States)
- South America (example: Brazil)
- Asia (example: China)
- Europe (example: the Netherlands)
- Other (example: Nigeria or Australia)
- Prefer not to say

4. Please select your highest level of education. If you are currently enrolled in a degree program, please select that degree. *

- Secondary Degree, High School, or Equivalent
- Bachelor's Degree or Equivalent
- Master's Degree or Equivalent
- Doctorate (PhD, MD) or Equivalent
- Prefer not to say


5. Do you consider yourself a native or fluent English speaker? *

- Yes - a native English speaker
- Yes - a fluent English speaker

6. Do you own or have you previously owned a smartphone? *

- Yes - I currently own a smartphone
- Yes - I do not own a smartphone now, but I have previously owned a smartphone

This content is neither created nor endorsed by Microsoft. The data you submit will be sent to the form owner.

 Microsoft Forms

Appendix

Appendix VI: Study Demographics


	Online Survey (Phase I)	Mock App Store Study (Phase II)	Semi-Structured Interviews (Phase III)
Age	62 Young Adults (18-24) 204 Adults (25-64) 7 Older Adults (65+)	25 Young Adults (18-24) 82 Adults (25-64) 2 Older Adults (65+) 2 Prefer not to say	1 Young Adult 17 Adults (25-64) 0 Older Adults (65+) 0 Prefer not to say
Gender	168 Women 95 Men 10 Other or prefer not to say	63 Women 42 Men 6 Other or prefer not to say	9 Women 9 Men 0 Other or prefer not to say
Education (have or in the process of obtaining)	214 Doctoral or Master's Degree 50 Bachelor's Degree 8 Secondary Degree 1 Prefer Not to Say	81 Doctoral or Master's Degree 19 Bachelor's Degree 7 Secondary Degree 4 Prefer Not to Say	15 Doctoral or Master's Degree 3 Bachelor's Degree 0 Secondary Degree 4 Prefer Not to Say
Nationality (by continent)	176 Europe 38 Asia 41 North America 18 Other or prefer not to say	67 Europe 21 Asia 13 North America 10 Other or prefer not to say	9 Europe 6 Asia 2 North America 1 Other or prefer not to say
Total	273 (147 Lose It! and 126 OpenLitterMap)	111	18

Appendix


Appendix VII: Heatmap of Significant Correlations for LoseIt! and OpenLitterMap

Spearman's rank correlation coefficients (ρ), where $p < 0.5$. *Value*, *Value App*, and privacy preferences for Lose It! (first); and *Value*, *Value App*, and privacy preferences for OpenLitterMap (second).

Lose It!



OpenLitterMap



Appendix

Appendix IIX: Question Bank for Semi-Structured Interviews

Opening Questions:

I ran a survey before this on values and privacy preferences back in the fall. Did you also do this [share screen with survey sample]?

If yes – how did you find rank values and privacy preferences on the survey? Can you walk me through your thought process?

If no – onto the Mock App Store.

Thinking back to the Mock App Store exercise [share screen with MAS]:

Which profile did you choose and how did you choose it?

What were you considering when downloading an application? Can you walk me through it?

How did you interact with notifications or pop-ups throughout the exercise?

Main Questions:

App Selection and Values

What is your process/what considerations do you have when choosing an app in everyday life? Could you walk me through it?

In what ways do you feel smartphone apps help or hinder you from reaching your goals? How does this affect your choice of apps?

How comfortable do you feel about the apps on your phone? Why?

Have you ever deleted an app on your phone? Why did you delete it? Did you download an alternative?

What app on your phone is your favorite – the one that feels the most “you”? Why?

App Selection and Data Privacy

What effect does the data an app collects about you have on your decision whether to download it? Why or why not?

How concerned are you about your privacy on your phone? How does this influence the apps you choose?

How do you feel about allowing apps and/or other parties – such as advertisers – access to your data? How does this affect which apps you decide to download?

Data Privacy and Values

Can you think of a time when something felt “off” about your privacy on your phone? What was the situation? What was it like for you?

How necessary do you believe it is for apps to collect data about you? Why?

How comfortable do you feel about the data collected on your phone? Why?

Are there particular kinds of data (such as your location, for example) that you feel comfortable/not comfortable with sharing? Why?

What concerns do you have about your privacy or the apps on your smartphone?

Perceived Level of Control/Frustration When Choosing Apps

How “in control” of your data privacy do you feel on your smartphone? How does this affect which apps you choose to download?

How much thought do you feel you give to your decision to download a smartphone app?

Have you ever felt like you “had” to have an app on your phone? If so, what were the circumstances? How did you feel about this?

What would be the best way to help you choose apps on your phone?

Closing Questions:

Do you have any additional questions?

Do you have any further comments on smartphone apps or data privacy? The floor is yours.

Thank you for your time.

Appendix

Appendix IX: Additional Quotations from Interviews

No.	Quote:
Q1	<p>Interviewer: [...] because of your disciplinary background, and because [...] you're already very in tune with values and [...] I guess, [the] theoretical backgrounds of these things. What would you consider would be the big values [...] when you're considering or looking at data privacy? Or when you're looking at downloading an app?</p> <p>P09: Um, that's a tough one. [...] I mean, again, it's all a cost-benefit analysis. So like, security does come in, but I - and privacy, which I guess are linked inextricably, [...] but I don't feel like they're [...] the ultimate value [...] I guess, it's utility [and] benefits [are] the most important thing really [...] [as] broad as that is, and also [pause] kind of a weird one, but like, not harming [anyone].</p>
Q2	<p>P26: [...] data protection is key in research. You need to protect every single person you interview, you need to make sure that whatever Dewey Decimal System [...] whatever system you [are] using [...] it's got to be so far removed from what that person really is.</p>
Q3	<p>P26: “[...] can we talk to the client?” “He's deaf, he cannot talk to you. That's why I'm phoning you.” “Okay. Now he needs to do X, Y, and Z.” Now, that's the way they can get it done. [...] Every government service should have an allocated, “this is the number you dial on WhatsApp. And you can FaceTime somebody at the bank.” [...] Somebody at social welfare, somebody at housing, somebody at the [tax office]. But they don't. [Exasperated laugh]</p>
Q4	<p>P08: [Privacy notices have] just become something people have to get through to get to the site they want and just tend to click the quickest buttons and the quickest buttons [tend] to be the ones that are highlighted and [tend] to be the ones that best give most data back to the company anyway. [...] I'm not a big believer in free will. And I think the way the manipulations happen around data, even when you have all the GDPR and stuff, and it's just [that] the forms and designs are just too clever for individuals to kind of completely overcome. We go online, we're looking for something, we want to get to it quick, and then this thing comes up, and you know, if you click “Accept All,” it's just going to bring you straight to the thing, where if you have to go to two more pages, and [...] click five more things. We click “Accept All.” [...] It's just the most common thing to do. [...] [As] Tristan Harris said, “the floor has been leaned that way,” [...] to make it easier [...] we're comfort-loving animals.</p>
Q5	<p>P15: One time when I opened a Facebook link in my Mozilla Firefox, when I went to [...] [my] browser history later [...] there was a log entry log for that [...] even though I don't have [the Facebook] app or something, even though I was not logging [into] it from my phone. [It was] from my Mozilla [...] I had [...] that entry in my browser history. I was super pissed that day. [...] I mean, honestly, that's so creepy. I think [...] they [linked the data to me] with my phone ID.</p>
Q6	<p>P22: I think it's personal information. [...] I'm using this to browse for my own purposes. I don't need ads [targeting] me. It kind of feels like [they have] a mind of their own. And I feel like decisions should be autonomous and not be influenced by whatever the ads [are trying] to feed me [...] I think, sometimes when we look through [...] a website or an app [...] for example [...] [say] I'm looking at this dress I really like, I could just be looking at it, because I could be thinking like, “oh, maybe it's a nice color.” And the thing is, that could be it [...] maybe I don't want to spend [money on this] but the thing is, the ad keeps popping up and saying that “hey, like...”</p> <p>Interviewer: “...here it is.”</p> <p>P22: [...] I don't need the extra prompt because [...] I only wanted that five minutes, when I was thinking about whether I should get the dress [...] [by] putting all [these] prompts in my life to try to make me, like sway, my decision [...]. I just [...] find it quite [disconcerting].</p>
Q7	<p>P06: So, I suppose it's the idea. I mean, so I think the common refrain, when that kind of thing happens, is people say, “Oh, they are listening. They're activating your microphone, and they're listening.” And that's unnerving enough. But that's [...] probably not it. What it probably is, of course, is they have formed such an effectively predictive algorithm of your behaviors, that they know what you want before you know you want it, or they know what you're going to discuss before you think of discussing it. And that's far more terrifying in my mind. So that's very unnerving. If they could just turn on your camera, well, that, you know, that'd be weird. But generally, I have, you know, “nothing to hide” [quotes done with fingers]. Not that it's okay for them to do it. But I'm concerned about the effects. However, the idea that they can predict your behavior ahead of time, that's very concerning.</p>
Q8	<p>P26: I guess with like your own personal development and [...] keeping your mental health and your wellbeing up, I guess, there's a certain point where you can nearly do too much. [...] there's so much out there now in terms of like the app store space that like pushes, “oh [...] this will help you” [...] for example, the productivity apps even they're like, “use this and you will be productive,” or “try this meditation app, and like, [...] you'll be extremely relaxed, and you won't be stressed anymore.” [...] you can only have so many meditation apps, like I already had one on my phone that I was using, when I saw [an] Instagram influencer, recommend another one [...] I don't need [...] two meditation apps [...] which are both taking information from me also [...] by “fall for that trap,” I suppose I just mean, there, because it's in line with my personal value[s], I think I need it. But actually, if I already have two apps that are doing the same thing, do I really</p>

Appendix

	need that? I guess with like your own personal development and you're keeping your mental health and your well-being up, I guess, there's a certain point where you can nearly do too much.
Q9	P15: Oh, another one irritating one is one, One Drive. I don't like the app at all. But I'm forced to take it because only because [university of employment] is promoting that app. And I feel in a lot of ways like it basically messes up with my work sometimes like it crashes [...] I would have preferred something like a Dropbox or something. But because of [university], I have to, I'm forced to download those apps.
Q10	P09: Like personal data spaces, I think that's a really interesting idea. And at the moment [...] they're being [...] presented as a private enterprise. And I think the idea of a government sponsored personal data space for its citizens, I think that'd be a really good idea. Rather than just having these sort of rules and guidelines about what corporations shouldn't do, and then they do it anyway. [...] I think that would be more effective. [...] Once the consumer can decide what terms they want to engage with the corporation with, with the other entity on [...] you can do it, or you cannot. But there's, you're not being forced either way, you're not being denied, let's say, what you might perceive as benefits of targeted advertising. And you're not being forced to engage with this you know [while laughing] shapeless, formless entity that's constantly soaking your data dry [joint laughter]. So potentially best of both worlds?
Q11	Interviewer: [...] you strike me as someone who's quite intellectually curious [...] you have <i>the Guardian</i> , Goodreads [...] you like [...] podcasts [...] how have you found apps to help you sort of in that aspect? If you don't find yourself intellectually curious, feel free to correct me, but [...] you strike me as someone [who is]. P14: No, I have, you know, I've really, some of those apps, I have to say, are really, really good. From that point of view, you know [...] certain ones, podcasts [...] or audiobooks now, and really interesting, as well. [...] Even Reddit and Twitter, from that point of view, depending on where your interests are focused. There is good stuff on there [...] it's not all bad [...] you're exposed to different things that you might normally not be [...] I've found those really enjoyable, you know, [<i>the Guardian</i>] as well. I enjoy their coverage.
Q12	P01: [...] another thing [...] I paid for YouTube [...] I hate watching ads on YouTube. [...] I don't know how much per month, five euros? [...] so, I can have YouTube and YouTube music on my phone as well that I can watch videos without, like with my screen off, and not get any ads on YouTube, things like that.
Q13	P01: The reason I care [about my data privacy] [...] is because [...] I do not trust companies enough to give them my data and trust them not to misuse it. For example, one company that does collect a lot of data, and I [I] do not care so much about, is Amazon. Amazon sends me emails every two days about new books coming out that I might like, and they know which books I like, because they have collected data on which books I bought before, which books I've read [...] on my Kindle, and so on. [...] But in this case, it provides me with something useful, and [...] I'm truly interested in the books that they are recommending. So even though they use the data [for] personal gains, I do not think that they are going to use it [...] with malevolence, or in a in a way that I'm not going to approve of. I am less trustful of companies like Facebook, for example, that have a record of misusing users' data for sort of at least gray purposes, if not downright evil purposes. [...] for example, when Facebook has collected data on users that has been used to target political statements, for example, to create echo chambers. YouTube is also bad in this regard. Because it's useful when [the AI] keeps recommending videos that I like. But the downside is that sometimes [...] it creates these echo chambers [...] when it gets political [...] I don't know if it keeps recommending me videos of Pokémon, because I've watched videos of people playing Pokémon, then I don't mind. But there is this rabbit hole where [...] once click on [...] some video or other and all of a sudden you have tons of videos of Neo-Nazis telling you things. And well, YouTube knows this [...] don't think it's necessarily their fault [...] because it just happens [with] the algorithm. I truly believe them. I just think that the algorithm tries to maximize your watching time by sending you videos that are outrageous [...] And once you get sucked into that, into that area, you keep watching those videos, or that class of videos. So, it's just a side effect of the algorithm. In the case of Facebook, I think it is even worse. I think that they have actually used the data to sell it to political agents that [...] wanted to bypass certain laws in some countries. [...] For example, in [home country], you have the day of the election or the day before the election, you're not supposed to show political ads, and so on, because it's one day where people have to think about who they're going to vote and so on. And things like Facebook [have] been used [...] by political agents to do [...] hidden campaigns. And it's always very subtle because it's not straight. It's not [a] downright campaign, but it's more highlighting certain articles, for example, blowing certain things out of proportion to certain a population to move them to vote. So, I think that's one way that Facebook has misused my data [...] I think it's one thing when you go to Amazon [...] because Amazon sells books. And sometimes books are about politics. So, what if they keep only recommending you books about the far right or the far left? [...] I never experienced that problem. In general, Amazon has a category that is politics. And they will recommend you books in that category of books, or science, or religion, or science fiction, or history. And you know what you are in Amazon for. You know, when you sign up that you're there to buy things. [...] You are not in Facebook to get political ads. You're in there for other things, and the fact that they use your data for something that is so different from what you originally signed for? I think it's, or the very least, you can say it's a bit perverse.
Q14	P06: [...] I'm very convinced that basically me as a person is not of interest to any [...] bigger company, me as a customer, maybe, or me as a consumer, maybe. So, I'm totally aware that [...] if I go onto Amazon, I get [...] linked products [...] that's okay for me. Because I still believe that this doesn't really impact my [...] real life in a negative way. I mean, I totally know that it would impact my life, because I get like a

Appendix

	recommendation for this book, or that book or whatever. But I don't feel like it would really negatively impact my life.
Q15	P15: One time I had an email in Gmail, I had someone sending me an email saying, "this is your password." And it was my one of my main password[s]. [...] I'm very careful about these things. Usually, I don't go to any kind of questionable websites or anything like that, [...] but somebody is sending me an email saying that "this is your email, so you better give me money." So those kind of things are also really scary.
Q16	Interviewer: What [does] secure look like for you? P23: [...] very simple, actually [...] as long as there [are] no unknown people contact[ing] me related to the information that I share. It's that simple, actually. Because sometimes, maybe you also experienced, that there is somebody calling you [and] you don't know who the person is [...] it's really bad for me. So, as long as the information that I share can be [secured?] by the people who need [it] [...] then it can make me feel comfortable for sharing it. I don't mind [sharing] the information.
Q17	P01: I do not think that Google has been found out as Facebook has [been], selling the data to political agents [...] Facebook with this investigation was truly found out to be selling political information about our people. And I think that can also be used for evil by governments if [it] wants to [...] but Google hasn't been found out to be doing this. [...] So, again, going back to how much I trust companies, I still trust Google a bit more than I trust, well, Meta. [...] I am less trustful of companies like Facebook, for example, that have a record of misusing users' data for sort of at least gray purposes, if not downright evil purposes.
Q18	Interviewer: So, I find it interesting [...] there's this [...] line that's being drawn here [...] between like data collection [...] by the companies, and then the role of government coming in too. And I wondering, how... I'm sensing a lot of trust with the government. And I'm curious to see, or to tease out a bit more, where that trust comes from, for you, in terms of the government as a regulator, and the government also as a data collector too? P09: Yeah, so I'm actually not usually big on governmental trust [joint laughter]. To be honest. [...] but this is a situation where it feels like it's the lesser of two evils in some ways. [...] Also, I think it depends on the government in question. I mean, I wouldn't have a huge amount of trust in China, collecting the data of its citizens, or, or, you know, [somewhat hesitating] with respect, even the US with, you know, Patriot Act, and... Interviewer: Totally ok! Why do you think I'm in this field? You can say... P09: [Joint laughter] Yeah? Fair enough. Fair enough. I mean, like, the [home country]'s government is usually terrible with data too. But that's usually due to gross incompetence, rather than malice. [...] so then, it's a matter of getting it right, as opposed to not doing it, if – oh, that's so glib! But you know what I mean? It's a matter of instituting an effective framework that minimizes incompetence. And in that [...], leaves [it] in the hands of people with power, who do at least theoretically have the citizens best interests at heart, as opposed to a corporate entity that is just about profit.
Q19	P23: Because instead of browsing by using [my] laptop or computer [...] I need time to do that [...] when I use my phone, I can just use it. Anytime I need it. A lot easier. More accessible, probably.
Q20	P02: I suppose [...] a lot of us [are] going towards sort of a more decentralized approach to sort of break down the ties. I suppose you can get a lot of these [...] if you used Tor browser, you can stop [...] the IP tracking. If you don't log in with your Google account. [...] you can try and separate these things out. [...] But the utility of having it all together is so huge. Am I prepared to give that up? I don't know. [...] Having things remember you is [in] inherent conflict with your privacy [...] but the utility it offers is pretty big. Is there a way? I don't know what the way of trying to protect it when it isn't essential, an inherently essentialized thing for that tracking? You could try and keep things local, but then [...] we don't really want things local, because then you're limited by your device.
Q21	P14: Facebook is awful social media [...] the functionality became [poor]. It was no longer [...] about the people you're hanging around with. You'd log on and 90% of the things on your profile were competitions or [...] stuff that you don't really want to see.
Q22	P13: [...] you have these apps for [...] this bicycle share in some cities, or the scooters and say, you know, you need it for this particular moment, but then you see in the description of things they have access to, and you're like, "wow." But [then] you're like, "Okay," [small laugh] "go ahead!"
Q23	P13: I think my defaults are just the necessary cookies. But I think [...] in some moments when it pops out, very often actually, it's not easy to, you still have to confirm. So, by default, it's just necessary cookies. But if I want to confirm the default, I have to click on a different button, that opens a different a full page about cookies. And then I can select my default. And otherwise, you have the very handy button that says, "accept all." It's either you know, you have "review your cookies," "accept all," or "reject." Sometimes I do reject, and then the page doesn't open anyway [...] Sometimes [...] maybe I'm not in a hurry to look at the page. So I go to the settings, and I select necessary cookies only. And sometimes I [am] just like, "Okay, go ahead."
Q24	P14: [...] [managing privacy notices is] a two second job [...] but we're like, "oh, [...] it's such a pain" [...] a lot of the time as well, I'd go through the settings, the app settings on the phone [...] go through the permissions [...] and change them [...] [you need to] pay attention to what they're looking for.
Q25	Interviewer: [...] what about those apps [...] do you like? P08: Mostly just convenience. [...] most apps I'm on is because I need to be on it [...], like WhatsApp. I'd have no friends if I didn't have WhatsApp. The [nationwide media organization app], I like it, because it's a non-curated feed. And it kind of acts more like the old newspapers, then then the social media apps, which gives you personalized feeds. So, we all get to see the same newspaper, you know. [...] I feel like the stories

Appendix

	<p>that are coming up on [nationwide media organization app] are coming up to everybody in the same order, pretty much. I don't do the comments, although some of them are funny, but [...] I try not to join them 'cause you'd waste half your life there as well. [...] I just feel like if I'm going on down the road, and I'm talking to somebody about something, they [have] seen on [nationwide media app], and see they saw the same story in the same fashion. And we'll both be talking about the same thing. Whereas if we're on social media, it's likely that it's been specifically curated because of its determination of my psychological profile. And it's really just there to grab my attention, specifically my attention. And I don't feel like that on [nationwide media app].</p>
Q26	<p>P30: [...] I think I'm relatively private in terms of what I put out, but like, they still know enough about me. And it could just be that, you know, it's like, because nothing ever goes away on the internet, right? [...] when the internet was new, I would have had more personal information on there [...] And that's enough information. So yes, it's very interesting. And like you're saying, because I have a marketing background, so I have worked on the other side of it. And I know how that targeting works. So, for example, I was helping a friend set up a business to do slumber parties for kids. And she set up a targeted ad for, you know, close to where she lives for a certain age range of males and females, because she was targeting the parents to try to get them to get slumber party's going for their kids. And she knows about the age range because of the same age as her own son. So, she knows they're around her age. And only the females really clicked on the links, so then she could create a new ad that just targeted the females in that age demographic and save money. And it's kind of it's like you're deliberately leaving the dads out of this opportunity to be a part of party planning for their kids because it's not lucrative. [...] So, I thought that was quite interesting. And sad, but you know, that's what we want, was because she's on a budget and trying to get a little business going. [...] But that's sort of [sad?] because yeah, maybe it was mostly women who clicked on it, but probably some dads clicked on it too. [...] I've only worked in marketing on sort of the small scale side of things so I haven't ever worked with a massive corporations that have access to loads of data or can buy loads of data. [...] I never worked in a place that I was sort of morally opposed to like in doing marketing for kids slumber parties, I'm like, "this is a good thing" [...] Whereas I think it can go a lot darker with other maybe larger companies that do buy chunks of data. And you know, they get your email data or your mobile phone data, and they bombard you with ads. [...] And that's something that I'm really opposed to. And I'd like to think that a lot of people who are coming into the marketing profession these days are also opposed to that. And they go more towards what's called "inbound marketing," where people come to you because they want to learn about you, rather than you going out to them. And just like constantly being in their personal space all of the time. But it depends. It's all about how the individual thinks about it, I guess.</p> <p>Interviewer: Yeah. No, thank you. That's [...] really interesting. And yet when you said about, like "getting it in your personal space," you're talking about like instances, like when you got the IVF stuff [you] mentioned earlier?</p> <p>P30: Yeah. [...] So that's sort of it. It's, there's, there's a line there for you that it's like, okay, slumber parties. I'm helping people, you know, like, give their kids a good experience. But then there's this. [...] IVF definitely could be helpful for a lot of people. And they might be welcoming that ad. I think what made me upset there was because I thought, what, what is it about me that's making this ad come to me? And like, what do they know about me? [...] When we did the targeting for the kids slumber parties, it was just based on age and location, so I could just as easily have gotten that ad. Um, it just somehow seems, [...] I don't know if I can really put into words what it is, it seems [...] less accusatory, like the IVF just made me feel like "[...] you should be doing something that you're not." [...] IVF is [a] very, very personal decision that they're targeting me on and making me consider.</p>
Q27	<p>P29: You can only have so many meditation apps. [...] I already had one on my phone that I was using when I saw [an] Instagram influencer recommend another one like I don't need two, two meditation apps, you know what I mean? Which are both taking information from me also. So yeah, from, by "fall for that trap," I suppose I just mean, there, because it's in line with my personal value, I think I need it. But actually, if I already have two apps that are doing the same thing, do I really need that?</p> <p>Interviewer: Yeah, yeah. And so, did you delete that other meditation app? Or did you keep both of them, out of curiosity?</p> <p>P29: [Laugh] Kept them both.</p>
Q28	<p>P08: YouTube is very handy, and very almost need to have, you know, just for looking up lectures, or just for general information, and of course, [...] the Google search engine. I mean, that gives us an amazing access to information [...] I don't think somebody could, could reasonably go about the world now without access to that information, without kind of maintainin' themselves as a stone age type person, you know? [...] I think [Google] gives us a "god complex," [...] that we think we know everything now. So, we don't need to learn. [...] We have Google in our pockets.</p>
Q29	<p>P14: [...] things like Twitter and Reddit, where you can spend hours just scrolling through, [...] they're good [for] knowledge about certain things, but [for] the amount of time you put into them, [they are] not necessarily worth the value that you're taking out of them.</p>
Q30	<p>P15: [The GDPR] helps a lot, actually [...] because I can feel the difference when I'm downloading or browsing [a] web page here and back in [birth country], because [...] I can see the list of cookies, right? Most of them at least. [...] from a user perspective, I have the power to disable [...] most of the cookies. But when I'm browsing the same page from [home country], there is no question, nothing. [...] I can assume that [...] these pages have cookies [...] and I don't have any option to disable [them].</p>

Appendix

Q31	P29: I think [the survey] definitely made me reflect on like, my own principles [...] I don't sit down and think about "okay [...] what principles are guiding my life?" Or [...] "what do I stand for in terms of [...] <i>Power?</i> " [...] I don't ever sit down and think about these. [...] When you're asked about them, it makes you reflect [...] on these things.
Q32	P07: Yeah, it was when the [survey] questions were about the app, [...] for example, this one [<i>Achievement and Hedonism Lose It!</i> questions on screen]. So [for] questions like this, I remember thinking, I do value these things. But [...] I just feel like I don't consider them right now when downloading anything, which was weird for me. So, I feel like even though I rated it, however I rated it, I don't feel like it truly matched [...] what I felt because of how I actually carry out my life online. [...] I almost feel like I wanted to say, "Well, okay, [...] it is important to me, you know, to do all these things, and I would consider it when downloading an app." But at the same time, I wouldn't let it stop me from downloading the app. So, I'm not sure if those if that was got across, you know what I mean?
Q33	P28: [I] kind of consider myself a Goal Setter in that [...] I live by writing lists, or I kind of think about that. I kind of quantify, to a point, my day. Yeah, but then I don't know – <i>Power? Hedonism?</i> Yeah, I didn't feel [...] [those aligned] with what I would consider to be my values. <i>Achievement</i> , I supposed, to an extent [...] I do feel I'm a Goal Setter. But I don't kind of think of it in a <i>Power</i> or <i>Hedonism</i> kind of a sense.
Q34	P22: [...] maybe I would say most days I could be an Adventurer, and then other days I could [...] [be] like a Helpful Neighbor. So that's why I found it hard to choose [...] based on my mood at that point, I could feel [...] I could be another one. So yeah, so I guess, it's not clear, it's not clear cut. But I think it's hard to be clear cut when it comes to like, personality.
Q35	P06: [...] I do understand that you prefer to feel secure over pursuing new experiences might hint at someone [...] who would also be very cautious with their data. And that maybe [...] the [app] store would guide me towards apps that have very strong data protection rules, but altogether, I wasn't.
Q36	P01: For example, I think I, I'm not sure if I did, but I thought that the one [...] Drink Water Reminder? One that [...] I picked, I think, for example, [...] it fulfills a need, and I liked the value, it provided me the value in the sense of I like the use that it provided me. While it didn't have [...] a green dot next to it. It had that one that didn't track too much. [...] I ranked the different apps first from the[ir] usefulness, so I will not pick something that I don't think I will use. [...] And then from those that [I] think could be helpful, I tried to pick the best [...] one that respected my privacy [best].
Q37	P14: And then, [...] getting the warning about these "apps aren't consistent with your profile" and then viewing the similar apps, but then sometimes the similar apps wouldn't have the functionality that you're looking for. So, you go back, and you download them anyway [...]
Q38	P13: I didn't see [the link to the Store]. [...] I [wouldn't] have thought to [push] on the link. Like, the brain automatically thinks, "Okay, thank you. You've done your work."
Q39	P08: Smartphones [...] are great for [...] making life easy, making life frictionless. But you know, we evolved in an environment through that friction, [...] our bodies and our psychology [...] our cognitive functions, and everything [...] has evolved through [...] frictions with nature, with life, with everything, with each other. And although convenience feels nice and gives us lots of dopamine [...] and frictionless life is going to, I feel, make for very fragile people. [...] That's a worry. [...] I think we need friction. [...] Friction [...] makes us. [...] It keeps us in touch with our environment. You know, when we rub against things, you realize that they're there. [The] frictionless lifestyle is dangerous for me. [...] there's a St. Andrews golf course in Scotland, they say it's the only golf course where the game was made for the course because that's where golf was invented. Every other golf course was made for the game of golf. I feel like the real world is where we were made for. [...] Everywhere online is just every other golf course that's been made specifically for us to make it frictionless and to make it soft, but it's not, it's not real. [...] we evolved bodies, we evolved legs, we evolved our brains to live in this world, in this nature. And the past 150 years we've lost touch with that, in the past 20 years we've [emphasizing] <i>really</i> lost touch of that, you know? [Joint laughter] [...] I think there's [...] very much a lot of dangers [...] [in] this frictionless style of life that we're pushing towards.

Appendix

Appendix X: Feedback Comments from Exit Survey

Comments were optional, with 25/66 exit survey participants leaving comments. Recorded here verbatim.

Question: What other features or information would help you select a smartphone application?

Feature	Participant Feedback
Suggesting Alternative Apps	<p>Comment #1: “I think features are the main driver of my decision. Sure privacy matters but not as much. If I could see a functionality based comparison alongside a privacy based, then privacy may factor in if I was compelled [and] the functionality I cared about was sufficiently similar. There were a lot of apps, navigating them was a substantial task, understanding the interface tools some time - collectively I feel these left me with little time to really think about apps and the privacy implications.”</p> <p>Comment #2: “To have the alternatives together with the pop up.”</p> <p>Comment #3: “Having apps most consistent with values listed first in App Store as there are so many apps available. Would like to have had more time to compare similar apps.”</p> <p>Comment #4: “Side-by-side comparison with similar apps.”</p> <p>Comment #5: “Being shown apps in the same category as the one selected ([e.g], see similar apps -> see other versions of that app type (e.g., period tracker) that align with values).”</p>
App Data Collection Information	<p>Comment #6: “Information on what linked/unlinked means (linked/unlinked to what?); filter by type of application and what is tracked - e.g. don't show applications that track identifiers; information on what it means for me, when they track a certain information (why is it bad for me?), maybe also how necessary it is for the app to track these information in order to work well.”</p> <p>Comment #7: “More details on what will get shared with others.”</p> <p>Comment #8: “Name, data collected.”</p> <p>Comment #9: “I think it would be good to know the definitions and repercussions of people having that information.”</p>
App Information	<p>Comment #10: “Whether the app is owned by a start up, a big conglomerate... If it is not free, whether the company that is part of a give back program, potential known ethical misconduct...”</p> <p>Comment #11: “Easy access to reviews.”</p> <p>Comment #12: “Reviews and ratings from others are helpful, however they can always be trusted. Brands sometimes post fake reviews and negative reviews can be overwritten when a new version of the app is released.”</p> <p>Comment #13: “Subscription period”</p> <p>Comment #14: “paid v unpaid, if you are an unpaid app does this lesson privacy”</p> <p>Comment #15: “More details about the app”</p> <p>Comment #16: “In app purchases”</p> <p>Comment #17: “Verified by Google play, comments, Popular friends with”</p> <p>Comment #18: “Usability, Accessibility, not having too many confusing buttons/clicks. Smooth interface. Good reviews”</p> <p>Comment #19: “Ability to store information gathered on myself directly to the Cloud or Dropbox, not to keep it on the phone,”</p>

Appendix

MAS User Interface	<p>Comment #20: “User-friendly, easy to use”</p> <p>Comment #21: “More directions as to what I am supposed to be doing - was clicking around but unclear about the aim”</p> <p>Comment #22: “App for learning Languages e.g. Duolingo and App for Wallpaper images”</p>
Other	<p>Comment #23: “Not all application choices may be directly related to personality profile of the user, it may also depend on the lifestyle followed due to other factors. Some some lifestyle choice questions, along with personality profiling can help better in choosing the apps.”</p> <p>Comment #24: “More info about the risks instead of just a general warning. I didn't really understand what each warning meant.”</p>

Appendix

Appendix XI: Reported Participant Rationale for Ignoring VcPA Selective Notices

Comments were optional. Recorded here verbatim.

App Being Downloaded	Comment
Impulse - Brain Training	Yes as it is something that interests me and might benefit my life
Mintal Tracker: Sleep Recorder	Because the app is needed, so I would put my need for the app above my need for security and confidentiality
Impulse - Brain Training	The app is a brain trainer, so even if it has trackers, it helps in stimulating my senses creating a sense of adventure
Onyx: Home Workout	Sometimes I may not get enough time to be outdoors, so a guide to simple workout is essential
Gratitude Journal Affirmations	I need to practice more gratitude!
Yoga for Beginners Mind+Body	I would like to know how to do yoga correctly.
SilverSneakers GO	I think it will help me even when I'm an adventurer
Bloom: CBT Therapy & Self-Care	It provides something different than the other apps (CBT)
Manifest - Affirmations	I was very interested in the Content and besides being a helpful neighbour, I indicated that I am not as concerned about data tracking
Manifest - Affirmations	I'll go through the app, just browse it, and maybe uninstall it later
Period Tracker My Cycle	I like to track my goals and every little detail. This looks like a useful app to me
Heart Rate Monitor - Pulse HR	More interested in gains from the app
Clover Period Tracker Calendar	I need it
Manifest - Affirmations	I want it
O My Soul Christian Meditation	Looks interesting
Mintal Tracker: Sleep Recorder	Looks useful
Heart Rate Monitor - Pulse HR	Looks useful
Bloom: CBT Therapy & Self-Care	I don't feel the relationship between the value profile and my choice means much. The app has utility I'd like, the privacy concerns are small
Heart Rate Monitor - Pulse HR	I like to track my heart rate as part of ensuring I am meeting my fitness goals

Appendix

Heart Rate Monitor - Pulse HR	I am not sure what tracked, linked, unlinked means, so I am not sure what they collect exactly. I'd like to use the app and the things collected don't look bad to me.
Impulse - Brain Training	The app appeals to me more than my concern over data collection
Bloom: CBT Therapy & Self-Care	Im interested in cbt as I study mental Health Nursing
Alive by Whitney Simmons	I know whitney simmons
Shine: Calm Anxiety & Stress	I've heard of this app before and wonder how it works
Mintal Tracker: Sleep Recorder	No alternatives
Drink Water Reminder	Tracker
Mantra - Daily Affirmations	I occasionally meditate
Impulse - Brain Training	I don't have too much personal info on my phone and surely they not hacking into my emails/what's apps?
Heart Rate Monitor - Pulse HR	Might be worth it for the data.
Mintal Tracker: Sleep Recorder	I don't really know the consequences of identifiers
Bloom: CBT Therapy & Self-Care	Have very good opinion on CBT as psychotherapy practice
Lunar - Period Tracker	Most useful to me practically
Paceline: Rewards for Exercise	I find I need something to help me/remind me to exercise regularly!

*And in the end
The love you take
Is equal to the love
You make*

The Beatles (“The End”)

The Ultimate Thesis Playlist (Privacy Belongs to Us)



[Apple Music](#)



[YouTube](#)