



University of Pennsylvania
ScholarlyCommons

Publicly Accessible Penn Dissertations

2016

Gay Data

Yoel Roth

Follow this and additional works at: <https://repository.upenn.edu/edissertations>

 Part of the [Communication Commons](#)

This paper is posted at ScholarlyCommons. <https://repository.upenn.edu/edissertations/1985>
For more information, please contact repository@pobox.upenn.edu.

Gay Data

Abstract

Since its launch in 2009, the geosocial networking service Grindr has become an increasingly mainstream and prominent part of gay culture, both in the United States and globally. Mobile applications like Grindr give users the ability to quickly and easily share information about themselves (in the form of text, numbers, and pictures), and connect with each other in real time on the basis of geographic proximity. I argue that these services constitute an important site for examining how bodies, identities, and communities are translated into data, as well as how data becomes a tool for forming, understanding, and managing personal relationships. Throughout this work, I articulate a model of networked interactivity that conceptualizes self-expression as an act determined by three sometimes overlapping, sometimes conflicting sets of affordances and constraints: (1) technocommercial structures of software and business; (2) cultural and subcultural norms, mores, histories, and standards of acceptable and expected conduct; and (3) sociopolitical tendencies that appear to be (but in fact are not) fixed technocommercial structures. In these discussions, Grindr serves both as a model of processes that apply to social networking more generally, as well as a particular study into how networked interactivity is complicated by the histories and particularities of Western gay culture. Over the course of this dissertation, I suggest ways in which users, policymakers, and developers can productively recognize the liveness, vitality, and durability of personal information in the design, implementation, and use of gay-targeted social networking services. Specifically, I argue that through a focus on (1) open-ended structures of interface design, (2) clear and transparent articulations of service policies, and the rationales behind them, and (3) approaches to user information that promote data sovereignty, designers, developers, and advocates can work to make social networking services, including Grindr, safer and more representative of their users throughout their data's lifecycle.

Degree Type

Dissertation

Degree Name

Doctor of Philosophy (PhD)

Graduate Group

Communication

First Advisor

Sharrona Pearl

Keywords

identity, LGBT, platforms, privacy, social media

Subject Categories

Communication

GAY DATA

Yoel Roth

A DISSERTATION

in

Communication

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2016

Supervisor of Dissertation:

Sharrona Pearl, Assistant Professor of Communication

Graduate Group Chairperson:

Joseph Turow, Robert Lewis Shayon Professor of Communication

Dissertation Committee:

John Jackson, Jr., Richard Perry University Professor
Joseph Turow, Robert Lewis Shayon Professor of Communication

GAY DATA

COPYRIGHT

2016

Yoel Roth

This work is licensed under the Creative Commons
Attribution-NonCommercial-ShareAlike 4.0 License.

To view a copy of this license, visit

<http://creativecommons.org/licenses/by-nc-sa/4.0/>

ACKNOWLEDGEMENTS

As the youngest of three siblings, I've always been able to follow in the impressive footsteps of my two older sisters. In high school, a physics teacher who had all three Roth children in his classes reacted to seeing my name on his roster by saying, "You've got some awfully big shoes to fill." Many years and countless hours of research later, this dissertation is the product of my continuing attempt to fill those shoes. First and foremost, my thanks go to Maayan and Nitzan, the best sisters anyone could ever ask for. Without your love, advice, commiseration, and persistent reminders to just finish writing my dissertation already, this would never have been possible.

Over the last four years, my advisor, Sharrona Pearl, has helped me cultivate my interests and curiosities into meaningful research. Her friendship and mentorship have made me a better researcher and teacher, and I'll forever be grateful for the time and effort she's invested in me. Joseph Turow and John Jackson have likewise nurtured this project and my interests, and I'm indebted to them for their questions, advice, and urgings to work outside of my comfort zone. I've also had the privilege of working with and learning from faculty outside my committee; many, many thanks to Carolyn Marvin, José van Dijck, Lance Wahlert, and Barbie Zelizer.

My decision to go to graduate school in the first place was shaped by the tremendous teachers I've had over the years. As an undergraduate, I was lucky enough to work with Simon Head, Paula Heinonen, Carol Nackenoff, Bob Rehak, Dominic Tierney, and Patty White. I'd also like to thank Litty Paxton for her mentorship in the art of teaching and managing an undergraduate class. Their dedication to undergraduate

learning, and the enthusiasm they bring to working with their students, has been an incredible inspiration.

At the Annenberg School for Communication, I've benefitted from a community of incredibly bright and engaged peers, whose ideas are present throughout this work. My thanks to Doug Allen, Chris Cimaglio, David Conrad, Nick Gilewicz, Corrina Laughlin, Deb Lui, Shane Mannis, Sara Mourad, Alexandra Sastre, and Aaron Shapiro for their support and friendship. The mentorship of Nora Draper has shaped this project and my own development as an academic in more ways than I can count; I owe her my immense gratitude for taking me under her wing. And, of course, my time at Annenberg wouldn't have been complete without my officemates Bo Mai and Sun-Ha Hong, whose companionship and conversation have made our office a truly special place to work.

To the many friends whose reflections on gay social networking apps I've shamelessly cribbed in this discussion: I hope I've done justice to your thoughts. Special thanks to Jane Abell, Jonathan Cowperthwait, Isaac Hock, Carter Green, Chris Kennedy, Ambar LaForgia, Arthur Nicholls, Laurie Voss, Shawn Walker, Kate Walton, and Natalie Zeldin. I'm grateful to all of you for being a part of my life.

Around the time that I first started researching gay social networking apps, I met someone on Scruff who would turn out to be both a source of academic inspiration and my co-pilot along the way. Nick Madsen has been a sounding board for my ideas, anxieties, frustrations, and excitements over the years, and has been patient and compassionate throughout even the darkest moments of this process. Thank you.

Last, but certainly not least, I've been unbelievably fortunate to have the unwavering support of my parents. Your pride in my accomplishments and your much-

needed nudges to get my work done have kept me going over the last four years. Despite my trepidation about having my parents read a manuscript that's in no small part about gay sex, I can't wait to share this with you.

ABSTRACT

GAY DATA

Yoel Roth

Sharrona Pearl

Since its launch in 2009, the geosocial networking service Grindr has become an increasingly mainstream and prominent part of gay culture, both in the United States and globally. Mobile applications like Grindr give users the ability to quickly and easily share information about themselves (in the form of text, numbers, and pictures), and connect with each other in real time on the basis of geographic proximity. I argue that these services constitute an important site for examining how bodies, identities, and communities are translated into data, as well as how data becomes a tool for forming, understanding, and managing personal relationships. Throughout this work, I articulate a model of networked interactivity that conceptualizes self-expression as an act determined by three sometimes overlapping, sometimes conflicting sets of affordances and constraints: (1) technocommercial structures of software and business; (2) cultural and subcultural norms, mores, histories, and standards of acceptable and expected conduct; and (3) sociopolitical tendencies that appear to be (but in fact are not) fixed technocommercial structures. In these discussions, Grindr serves both as a model of processes that apply to social networking more generally, as well as a particular study into how networked interactivity is complicated by the histories and particularities of Western gay culture. Over the course of this dissertation, I suggest ways in which users, policymakers, and developers can productively recognize the liveness, vitality, and durability of personal information in the design, implementation, and use of gay-targeted

social networking services. Specifically, I argue that through a focus on (1) open-ended structures of interface design, (2) clear and transparent articulations of service policies, and the rationales behind them, and (3) approaches to user information that promote data sovereignty, designers, developers, and advocates can work to make social networking services, including Grindr, safer and more representative of their users throughout their data's lifecycle.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	vi
Table of Contents	viii
List of Illustrations	ix
Introduction, Literature Review, and Methods	1
<i>Data</i>	12
<i>Gay</i>	35
<i>Method</i>	48
Birth	70
<i>The infrastructural layer</i>	73
<i>The personal layer</i>	83
<i>The social-spatial layer</i>	98
<i>Conclusions: Asymptotically approaching embodiment</i>	108
Life	114
<i>Objectionable, indecent, and pornographic</i>	118
<i>Normative platforms</i>	130
<i>Negotiating gay visibility</i>	133
<i>Expressive resistance</i>	139
<i>Conclusions: The least restrictive alternative?</i>	144
Afterlife	149
<i>Data's risky afterlife</i>	157
<i>Platform and context specific data flows: A model</i>	162
<i>Commercial afterlives</i>	165
<i>No fats, no femmes, no privacy?</i>	185
<i>Conclusions: Making sense of revenge porn</i>	216
Conclusion	221
Coda: Death?	233
<i>Symbolic death</i>	235
<i>Real death</i>	244
Works cited	254

LIST OF ILLUSTRATIONS

Figure 1: Grindr profile creation screen	88
Figure 2: Completed Grindr profile	95
Figure 3: Grindr Cascade	101
Figure 4: Filters for Grindr Cascade	103
Figure 5: Grindr content review placeholder image	123
Figure 6: Scruff content policy notices	126
Figure 7: Public view of “Wanna Play?” installation	149
Figure 8: Alternate public view of “Wanna Play?” installation	152
Figure 9: Data flow model	163
Figure 10: Marketing on Grindr promotion	179
Figure 11: Recent tags cloud from Douchebags of Grindr	198

CHAPTER 1

INTRODUCTION, LITERATURE REVIEW, AND METHODS

Gay social media is moving out of the margins and into the popular spotlight. Since its launch in 2009, the geosocial¹ networking application Grindr has garnered consistent attention in the mainstream press. Vanity Fair referred to Grindr as “the world’s biggest, scariest gay bar” (Kapp, 2011). One of the earliest mainstream stories about the app, published in *The Guardian* in 2010, stated, unequivocally, that Grindr “is reconfiguring the landscape of human relationships” (Vernon, 2010) — a lofty achievement for any smartphone application, much less one that had been on the market for less than a year. In a recent feature about the app published in *The New York Times*, the author confirmed Grindr’s status as “the killer networking app in gay social media” — noting that the service has inspired droves of imitators, seemingly boundless criticism from users and pundits alike, and continued, exuberant engagement by millions of users worldwide (Trebay, 2014). And, just days before the 2015 Super Bowl, Grindr competitor Scruff bought a 48-foot billboard in the University of Phoenix Stadium parking lot, portraying two men in a locker room with the caption, “Play on our team” — a marketing move that garnered national headlines and, according to a spokesperson for the service, resulted in a 20 percent increase in profile creations in the Phoenix area (Mosendz, 2015).

As services like Grindr and Scruff become increasingly popular and pervasive — as the glow of smartphone screens becomes a phenomenological mainstay of gay bars, and of gay life more generally — we’re faced with the task of unpacking the

¹ Geosocial media are social networking services or platforms that use geolocation data (such as GPS coordinates) to connect users with each other on the basis of geographic proximity.

consequences of these new electronic faces of gay sociality. In this dissertation, I examine a fundamental process at the heart of these services: the transformation of the gay body into data. Through increasingly sophisticated and popular pieces of software, the body's dimensions, contours, and qualities — and even its position in geographic space — are rendered as items in a database. The multifaceted, complex lived experiences of gay male identity and sexuality are translated into an assemblage of data points that can be aggregated, acted upon, managed, and outputted in a variety of forms, to a variety of different ends. The central concern of this dissertation is exploring the mechanisms by which this process takes place, and identifying how this practice of making-into-data influences the possibilities for self-expression and safety for gay men in an increasingly mediated environment for interpersonal interaction.

The depth and breadth of these practices of data-driven self-expression have evolved over more than three decades of popular use of networked technologies by gay men. Each generation of gay media — from pay-per-use text messages on the pre-internet French networking platform Minitel to apps taking advantage of always-on broadband mobile data connections — has enabled certain types of embodied self-expression, while constraining others. While it's easy to assume that self-expression has become more straightforward, less constrained, and safer in line with improvements in the underlying technologies that enable that expression, this project examines the complications and concerns that accompany the most recent generation of gay services. Many of these concerns — about expression, authenticity, and commercialization — have existed throughout the more than three decade-long history of gay networked media; but some,

such as how user expression is managed in complex commercial and regulatory frameworks, constitute new and relatively unexamined developments.

Geosocial media like Grindr represent an especially important site for examining these processes. Using new technologies like the integration of Global Positioning System (GPS) capabilities into mobile phones, these services enliven user profiles with information that would have been difficult if not impossible to share online in previous decades. And, unlike the previous generation of browser-bound gay social networking services (such as Manhunt, Gaydar, and Gay.com), mobile apps like Grindr make gay sociality accessible almost anywhere a cellular data or WiFi connection is available. The result has been both an increased depth of user expression, alongside a significant broadening of the appeal and user base for gay-targeted social networking services.²

Of course, this increased depth of online self-expression is not unique to gay men or gay-targeted social networking services. A wide range of recent applications, services, devices, and technological systems have allowed people to collect, evaluate, and share information about their bodies and personalities in increasingly granular ways. Steve Mann, an early pioneer in wearable computing, has chronicled his experiences over more than three decades as a practitioner of what he calls “sousveillance”: the use of body-borne monitoring technologies for individuals to gather enormous amounts of audiovisual information to better understand themselves and their environments (Mann, 2005; Mann, Nolan, & Wellman, 2003). The Quantified Self movement takes wearable sensors, as

² While virtually all gay-targeted social networking services claim to have “millions” of users, few offer more specific information about the number who are actually active each month. In one of the few detailed accounts of active user information, Grindr’s July 2013 press fact sheet notes that the service has “more than 7 million users in 192 countries,” with more than 300,000 users logged in at any given moment. We can infer from Grindr’s disclosure of this data (a practice not adopted by their competitors) that their active user count compares favorably to those of competing services.

well as other self-monitoring devices, as an opportunity to make “big data” personal, transforming large data sets into tools of personal empowerment — an argument that has had particular resonance within the domain of healthcare (Nafus & Sherman, 2014; Swan, 2012). Popular wearable devices like Fitbit and Nike+ pedometers have made motion tracking an increasingly common practice amongst fitness-minded individuals. And, within the arena of interpersonal interaction, an entire industry of data-driven matchmaking services (like Match.com and OkCupid) employ detailed surveys and extensive profile creation forms to algorithmically match users with each other. In all of these cases, we find a significant expansion of the types and quantity of information collected through devices and software as a means to understanding human behavior and interaction.

More generally, the massive popularity of mainstream social networking sites like Facebook have made these practices of data-driven self-expression through software seem mundane and everyday to millions of people worldwide (Baym, 2010; Marwick, 2005; van Dijck, 2013c). Across a wide range of networked platforms, people share details about their bodies, identities, interests, interpersonal networks, business aspirations, and quotidian goings-on in ways that make human activity legible to and through software. And, increasingly, we invest an enormous amount of time, energy, and emotion in our online presences. We no longer imagine our social media profiles as avatars, to be invented and consumed within the confines of cyberspace; rather, our bodies, identities, relationships, and activities are deeply intertwined with our presences on networked platforms.

In this work, I'm interested in tracing the origins and existences of these intertwined presences; how we invest ourselves and our identities in them; how they present a version of us to strangers, friends, and advertisers; who we entrust with watching over them; and what we can do to make the sites of networked self-expression safer and more representative of the diverse values and identities of the users and communities who constitute them. In order to examine these questions, my research critically studies how the design, management, regulation, and everyday use of Grindr, the most popular gay-targeted geosocial networking service and the pioneer in the industry, impacts the bodies, identities, and communities that gay men are able to express and form.

This project examines the interplay of bodies, identities, and digital information as part of a process of networked self-expression on gay social networking applications. Throughout this work, I argue in favor of a model of networked interactivity that conceptualizes self-expression as an act determined by three sometimes overlapping, sometimes conflicting sets of affordances and constraints: (1) technocommercial structures of software and business; (2) cultural and subcultural norms, mores, histories, and standards of acceptable and expected conduct; and (3) sociopolitical tendencies that both are represented as and are popularly understood to be fixed technocommercial structures. This tripartite framework explicitly rejects the two dominant paradigms for conceptualizing identity formation online: first, the suggestion that networked identity construction is an autonomous, individual act of self-authorship; and second, a parallel argument that identity construction is overdetermined by technocommercial structures that are indifferent to the particularities of cultures and communities. While situationally

insightful, I argue that neither position can adequately account for the on-the-ground creation, management, circulation, and use of personal information and networked identity. Across the following three chapters, I use the gay-targeted geosocial networking application Grindr as a case study for the model of networked self-expression I develop. Grindr serves both as a small-scale model of processes that apply to social networking more generally, as well as a particular study into how networked interactivity is complicated by the histories and particularities of Western gay culture.

Across the following three chapters, this work is organized around six recurring primary research questions:

- How do the interfaces, policies, and practices of gay-targeted geosocial networking services influence the types of bodies and identities that users are able to express?
- What structures of interaction, observation, and expression are enabled by the novel interfaces and implementations of geosocial networking services? How do these practices of looking and interacting relate to offline practices of gay sociability (e.g. cruising)?
- How does the management of user-generated content on gay-targeted social networking services impact the visibility of non-normative bodies and identities?
- How are the structures of governance of gay-targeted social networking services — e.g. policies, moderation practices, etc — presented to users? What is the public response — from users, journalists, bloggers, etc — to these policies?

- How — and by whom — is the information about users gathered by and displayed on gay-targeted social networking services used? How, if at all, are these different uses of personal information made visible to users?
- How can software interfaces and service policies be crafted or refined to enable diverse expression, allow for innovative user behavior, and protect individual and group privacy and safety (both online and offline)?

These questions are ordered to build on each other to construct a ground-up picture of the structure, operation, and use of gay-targeted social networking services. They are also organized to reflect what I term the “lifecycle” of gay data: its creation, primary use, and circulation.

Throughout this work, I use the idea of a lifecycle as a structuring metaphor for the different practices of data creation and use I examine. The term “data lifecycle” has previously been used within the domain of information security to indicate concerns about how to protect the accessibility and security of data over extended periods of time — a question of archival and formatting strategies that is primarily of concern to individuals working within corporate information technology settings. But, more recently, some academic research in human-computer interaction (HCI) has adopted lifespans as a way to conceptualize the persistence of information systems across multiple generations of people — a vital recognition that data created today can remain critically important once its creators are no longer present to use it (Friedman & Nathan, 2010; Friedman, Nathan, Lake, & Grey, 2010; Nathan, Lake, Grey, & Nilsen, 2011; Yoo, Lake, Nilsen, Utter, & Alsdorf, 2013). These studies treat information not solely as an entity in

the present, but rather as a sociotechnical construct whose uses and impacts can persist and change over time.

By using the idea of a lifecycle as a way to conceptualize the practices of data-driven self-expression I examine, I want to stress this tendency toward change over time. Each chapter of this dissertation takes up a particular moment in the “life” of gay data: from its creation through the interrelationship of users, devices, and software interfaces, to the myriad unpredictable applications of personal information that emerge outside of data’s immediate, intended uses. Over the course of each chapter, I suggest ways in which users, policymakers, and developers can productively recognize the liveness, vitality, and durability of personal information in the design, implementation, and use of gay-targeted social networking services, with an eye toward making these services safe and representative of their users throughout their data’s lifecycle.

Chapter 2, titled “Birth,” takes up the processes and practices by which the gay body is translated into data. This chapter focuses on the devices and software interfaces that structure this act of translation, and investigates the ways in which particular software design conventions enable certain types of expression and interaction, while systematically constraining others. Using the technique of interface analysis, I critically examine the interface and interactive experience of the Grindr application, identifying how the Grindr profile creation process both puts users under external surveillance, and encourages them to carefully surveil themselves. I argue that the Grindr application is a space of *vertical mediation*, in which multiple types and layers of data converge within a single interactive scene to create new, data-driven subjectivities. These subjectivities are the result of technological practices of observation and disclosure, and represent a new,

machine-legible form of gay identity. The subjects represented in the Grindr application are at once vital, messy, and alive, even as they are deeply embedded in machinic structures of automated sorting, filtering, searching, and processing. While I recognize the constraints of a small screen and the need for simplicity and ease of use, I recommend less structured interface designs that give users the capacity to freely express information beyond the bounds of drop-down menus.

Chapter 3, titled “Life,” interrogates how gay data is managed and controlled. This chapter focuses on the practices of content management that structure what users are allowed to share about themselves on gay-targeted social networking services. While managing user-generated content is a common practice across social networking services, the policies implemented on gay-targeted services tend to be distinctively restrictive in scope and highly specific in formulation. I identify the technical, legal, and social affordances that authorized the creation of these policies, as well as the consequences for their implementation on user behavior. I locate the policies and practices of gay social networking services within broader discussions of acceptability, mainstreaming, and proper self-expression within gay male communities, and argue that service providers (including the developers of Grindr) intervene in these normative contestations in a way that promotes a banal, minimally erotic version of gay visibility in networked contexts. Through an examination of the vernacular practices of expression that users deploy to negotiate the constraints introduced by formal content management policies, I suggest that informal resistance is an essential part of how users engage with the imposed constraints of social network policies.

Chapter 4, titled “Afterlife,” takes up uses of gay men’s personal information that go beyond the manifest functions of gay-targeted social networking services. I examine two cases that illustrate the potential consequences of unexpected or unauthorized uses of personal information originating on gay-targeted social networking services: (1) In-app advertising and the “freemium” commercial structure of these applications; and (2) The blog Douchebags of Grindr, which publicly posts screenshots of Grindr profiles that its authors deem offensive or inappropriate. In both cases, personal information is used by an agent other than the data’s original creator, for purposes that differ from the original goal of enabling networked social interactions between gay men. But, critically, the structural characteristics of these data flows differ in fundamental ways — as do their corresponding risks. Drawing on social, legal, and technical remedies, I argue that each class of data flow requires a differently-tailored, but ultimately generalizable, solution. While all flows of personal information across social contexts and technical platforms introduce some risks to user safety, I argue for an approach to data management that prioritizes transparency and personal data sovereignty while minimizing constraints on the ability of users to deploy networked media in innovative or unexpected ways.

Finally, I want to emphasize that the questions and discussions outlined above are intended to speak both to gay-targeted social networking as a genre of social software with its own particular conventions, histories, and complications, and to gay-targeted social networking as a representative case study of broader practices and trends in the creation, management, and circulation of identity and information in networked contexts. The title of this work, *Gay Data*, stresses that there is something uniquely gay about the data in question; and, in many instances, this is in fact the case. I argue that the history of

gay male culture in the United States — and particularly, gay men’s ongoing contestation of group identity and public self-expression — has resulted in patterns of software design and development that are unique to gay men, and are worthy of study in their particularity. But I also want to keep the broader questions of personal information and identity — the more generic meaning of “data” in *Gay Data* — within the scope of each of these chapters. As critics and practitioners, we should look to gay-targeted social networking as a way to see the emergence and evolution of new conventions of online interaction, often years before their rise to mainstream popularity. The recommendations I outline for the design or governance of Grindr can, in many cases, be readily extended to other platforms — and I draw explicit connections to the designs, policies, and practices of mainstream platforms like Facebook, Instagram, Tinder, and OkCupid (and others) throughout this work. The scale of these services may differ from Grindr — billions of users on Facebook, to Grindr’s millions — but the underlying questions of inclusiveness, expressiveness, and safety by design do not differ substantially. The experiences of gay men on a niche social networking service speak to the experiences of multiple different constituencies of users on mainstream platforms. Where possible, I stress these points of productive overlap, and try to identify how we can use a study of niche networks as a model for networking more generally.

The literature review offered in this chapter is organized to mirror this polyvalent approach to the gayness of *Gay Data*, as well as the broader questions of data that form its backdrop. I begin by tracing academic research into identity, online selfhood, and commercial social networking at a broad level. Based on this discussion, I then offer a specific examination of gay culture and self-expression, both online and off. These two

strands, held momentarily apart but ultimately intertwined throughout this dissertation, form the backdrop for the discussions that follow.

Data

Atoms, bits, and the technosocial self

In what has now become a classic dualism, Nicholas Negroponte's *Being Digital* (1995) gave language to one of the central theoretical concerns of the so-called "Information Age": the separation between atoms and bits. Negroponte predicted that, in short order, many of the physical objects (that is, atoms) that we interact with on a daily basis would be replaced by digital copies (that is, bits). This shift, chronicled in *Being Digital* with reference to a wide range of technologies — CD-ROMs, e-books, letters, and so on — constituted a critical articulation of a development that a wide range of academic researchers and popular observers had tried to make sense of: the vexing intermingling of the physical world with the ephemerality of digital content and interactions. In the ensuing two decades of empirical and theoretical research into information technology, a wide range of researchers have continued to engage with precisely this question — from political-economic analyses such as Manuel Castells's *Rise of the Network Society* (2000) to Lev Manovich's humanistic *Language of New Media* (2001).

As a way to articulate the theoretical stakes for this discussion, I want to take up a subset of these questions. The tensions articulated by Negroponte, Castells, Manovich, and others — sometimes framed as a contestation of "digital dualism" (Jurgenson, 2012) — are similarly reflected in long-standing theoretical debates about the role of the body (atoms) in computer-mediated interactions (bits). Beginning with some of the earliest

accounts of networked identity formation and interaction on bulletin boards and text-based chat rooms, I trace these debates as a way to situate this study's treatment of bodies and technologies as fundamentally intertwined. By mapping the progression of these theories, from utopian conceptions of possibilities of the invention of digital selves to concerns over the consequences of different modes of online self-expression, this section outlines a theoretical middle ground: a version of online embodiment that recognizes the necessary relationship between the digital and the physical self, while privileging neither.

A key entry point into the narrative of digital embodiment is Julian Dibbell's "A Rape In Cyberspace" (Dibbell, 1994). The facts of Dibbell's narrative are well-known: Within the universe of the online roleplaying community known as LambdaMOO, a character going by the name of Mr. Bungle used a piece of software mimicking a voodoo doll to force other players to commit sexual acts in a public space in the game against their will. The details of the acts were spelled out in highly graphic terms, and were immediately understood by members of the LambdaMOO community to constitute rape. The individuals whose characters were targeted by Mr. Bungle couldn't react to the actions taken against them, or stop the process; they could only sit back and watch in horror as their online selves — or some version of their online selves — were brutally violated in what, moments prior, had been a tightly-knit and vibrant networked community. One of the individuals involved in what Dibbell termed "The Bungle Affair" later confessed to him that, during the ensuing public debate on LambdaMOO about the consequences of the rape, she frequently found herself in tears at her keyboard. She told Dibbell that she was physically shaken when she was forced to recall the lines of text describing the violation of her avatar that she read on her computer's screen during the

rape. The tension at the heart of “A Rape In Cyberspace” is the dissonance that, despite the obvious psychological gravity of the situation for members of the LambdaMOO community, “no rape at all as any RL court of law has yet defined it” took place (Dibbell, 1994, p. 473). What does “rape” mean when its mediated context is presumed to be fundamentally characterized by abstraction and ephemerality?

Dibbell rejects this presumption of ephemerality and disembodiment. He writes that in LambdaMOO, as in other social spaces,

Amid flurries of even the most cursorily described caresses, sighs, and penetrations, the glands do engage, and often as throbbingly as they would in a real-life assignation — sometimes even more so, given the combined power of anonymity and textual suggestiveness to unshackle deep-seated fantasies. (Dibbell, 1994, p. 476)

This stands in stark contrast to traditional images of computer users as seated at a keyboard, disengaged from the lives unfolding on the screen. In many early accounts of networked interactivity, bodies were either nonexistent in online spaces, or were fantastical simulations of their offline counterparts. They existed, in Juniper Wiley’s terms, as part of a “a parallel universe” (Wiley, 1995, p. 161). Similar positions about the operation of identities online have been put forth, canonically, by Donna Haraway (1991), Sherry Turkle (1995), N. Katherine Hayles (1999) and Howard Rheingold (2000). Sherry Turkle, in *Life on the Screen* (1995), described online identity work as “freedom” from the burdens of physical bodies, enabling people to become better versions of themselves. Howard Rheingold’s account of The WELL in *The Virtual Community* (2000) chronicled the equalizing power of networked contexts, noting that social and professional positions are invisible and do not directly grant privileges to those who possess them — making The WELL egalitarian in a way that the offline world could

never be. Allucquère Rosanne Stone (1991) characterized online identities as aligned with hybrid, Mestiza subjectivities, free from the constraining dichotomies and impositions of offline social systems. And, in *My Tiny Life* (1991), the book-length account of his experiences on LambdaMOO, Dibbell describes with enthusiasm his creation of the persona of “Samantha,” an online exploration of the life as a woman that was otherwise inaccessible to Dibbell as a biological male.

In each of these cases, we see what Jenny Sundén (2003) characterizes as people “typing themselves into being”: a conscious act of self-authorship enabled by the particular affordances of networked media. This, ultimately, is the phenomenon at the heart of the “Bungle Affair.” Mr. Bungle was, in actuality, a group of undergraduates at NYU, sharing a single online persona. Their malicious actions were a typing-into-being of an act of violation whose emotional dimensions were readily apparent to LambdaMOO’s other participants. By authoring the texts of Mr. Bungle’s persona and the act of rape, a group of undergraduates were able to enact those activities within a networked context where writing was understood to be equivalent to being. Dibbell’s position in “A Rape In Cyberspace” argues explicitly against a view of cyberculture that privileges a disembodied utopia of virtual identity construction freed from offline constraints or consequences. LambdaMOO users do not merely perform their characters as independent identities; rather, they are intimately connected to them, and experience significant real-world consequences as a result of online actions. The broader theoretical claim — that “RL” bodies should not be imagined as independent of their digital manifestations — has emerged as a dominant trope in the academic literature of cyberculture. I take this middle ground — an assertion of the fundamental

interconnection between “RL” bodies and identities and networked activity — as the theoretical premise on which my analysis of gay male identity and activity on social networking services rests.

Despite the possibilities for new forms of expression offered by networked media, it’s important to recognize that particular technologies also bring with them particular sets of constraints on the ability of individuals to share information about their bodies and identities. Especially striking in early studies of computer-mediated-communication was the issue of bandwidth: the amount of data that can be exchanged in an interaction with a service. Text-based chat protocols, like LambdaMOO and Internet Relay Chat (IRC), were built for the slow and high-latency internet connections of their time: users exchange relatively small amounts of data while connected to IRC. This has important consequences for user experience: As Stone (1995) notes, compressing higher-bandwidth embodied interactions into lower-bandwidth forms (for example, translating the embodied experience of sexuality into phone sex or online chats), creates significant potential for loss or distortion.

The emergence of newer technologies — most notably, high-speed internet connections — augmented this process, reducing the need for and negative consequences of compression (Shaw, 1997; Waskul, 2002). Higher-bandwidth exchanges become what Jason Farman (2012) calls “sensory-inscribed”: the body as a feeling entity is bound into mediated processes of communication. Nevertheless, the reduced need for compression should not be mistaken for an equivalence between online and offline selves; Farman argues that, phenomenologically speaking, “full, embodied presence is always being deferred” (Farman, 2012, p. 30) — that even those interactions we imagine to be highest-

bandwidth, such as a face-to-face conversation, always involve a series of textual, visual, and embodied significations that defer “direct” engagement. Users interacting online traffic in the symbolic currency of particular networked spaces, reflecting the affordances and constraints of that social context, as would individuals speaking face-to-face — a position reminiscent of both classic sociological performance theory (Goffman, 1959) and the encoding/decoding model of communication (S. Hall, 1973).

Even where networked media might be able to accommodate (in terms of bandwidth) a robust presentation of self, they are not always successful in doing so. In practice, design choices about how users disclose information about themselves can often force reductive self-presentations that bear little in common with the complexities of raced, gendered, or embodied experiences of self. As software tries to create order among the enormous set of possible identity presentations by giving users a series of clear choices, it may have the effect of constraining the identities users are actually able to construct. As Lisa Nakamura puts it,

[Interface features enforce] a menu-driven sense of personal identity that works by progressively narrowing the choices of subject positions available to the user, an outcome that seems to fly in the face of claims that the Internet allows for a more fluid, free, unbounded sense of identity than had been available in other media — or, indeed, in the world — before. (Nakamura, 2002, p. 104)

These menu-driven identities stand in stark contrast to the “unclickable, hyphenated, hybrid, ‘messy’” forms of identity that are essential in reflecting embodied experience (Nakamura, 2002, p. 120). Users may feel more comfortable approaching interfaces with fewer choices but significant detail is lost in the process. Some of the design choices underpinning these interfaces are the result of technological necessity, a reiteration of the

classic bandwidth problem; but others reflect normative judgements about how bodies — and, specifically, bodies that are potentially marked as non-normative — ought to be expressed online. This dissertation examines the effects of this nexus between technologies, social norms, and design practices in the everyday experience of using gay-targeted social networks, with the intention of highlighting how the design, implementation, and management of these services consistently constrains the expression of certain bodies, identities, and practices, while permitting the disclosure and public display of others.

Identity authorship in networked contexts

Taking a step back, I want to explicitly identify the theoretical underpinnings on which these accounts of identity construction are based. Specifically, across a wide range of scholarship focusing on networked identity, the notion of identities as performances has become an important way to describe how individuals create and understand their presentations of self in online social spaces.

The seminal account of this performative model of identity emerges in the work of sociologist Erving Goffman (1959). Individuals in social spaces, Goffman suggests, are akin to actors on a stage. Actors — that is, all people — constantly engage in what Goffman calls “impression management” by carefully controlling the information they reveal about themselves in their interactions with others. Based on their understandings of different social situations, individuals make choices about what information to reveal — and how to reveal it — in order to present themselves in the best possible light. These choices, Goffman notes, take place within institutional and normative contexts that

influence the performances of the individuals they govern. A sensitive and successful presentation of self reacts to information available to individuals both about their social partners, and about the broader context within which an interaction takes place. These presentations, Goffman suggests, are akin to the behaviors of actors on a stage in a drama, carefully manipulating the impressions of their audiences.

Networked settings seem like an ideal instantiation of Goffman's dramaturgical model, even as they complicate some of its essential components. On one hand, the permissiveness of many networked settings appears to give individuals the freedom to perform their identities as they wish, constructing performances of self that are liberated from the embodied constraints of physical co-presence and the limitations of one's physical appearance and manner. Individuals have the ability to selectively reveal information about themselves based on their understandings of the purpose or norms of a given context (boyd, 2007; Papacharissi, 2002a; 2002b; 2009; 2012; van Dijck, 2013c), or expectations they have of their audiences (Marwick & boyd, 2011). Of course, as Nakamura and others recognize, this is an ideal, rather than a reality; the limitations of software interfaces constrain identity performances. Nevertheless, the relative anonymity afforded by many online contexts, coupled with the emphasis on selective disclosure embedded in the logic of profile creation, gives individuals the ability to purposefully construct identity performances that reflect the selves they wish to share with others.

Significantly, however, this permissiveness also diminishes the cues available to individuals about what the properties of a given context actually are, and what an appropriate presentation of self might look like. Whereas a Goffmanian drama takes place within the bounded spaces of an interpersonal stage, within which an operating

consensus can readily be established by the participants as to what kinds of performances are allowed, online settings are considerably more ambiguous. Context collapse — a condition in which multiple social situations commingle, potentially unexpectedly and not necessarily in line with the wishes of the participants in those situations — highlights the fragility of the frames within which online identity performances take place (Davis & Jurgenson, 2014; Duguay, 2014; Marwick & boyd, 2011; Wesch, 2010).

Commercial structures of self-expression

Of critical importance in these discussions are the ways in which software developers and designers strategically craft social services to make individual self-expression and interpersonal interaction into a commercially-lucrative product. This takes place on two levels: (1) By linking expression and marketing in a way that transforms individual activities into a form of free labor; and (2) in the reflection of commercial logic in the marketing, management, and design of apps and websites.

Tiziana Terranova's discussion of free labor anchors the vibrant forms of cultural production that are characteristic of many narratives about individual self-expression on the internet within the corporate systems of labor and value that unavoidably structure the operation of the web.

Within the early virtual communities, we are told, labour was really free: the labour of building a community was not compensated by great financial rewards (it was therefore 'free', unpaid), but it was also willingly conceded in exchange for the pleasures of communication and exchange (it was therefore 'free', pleasurable, not-imposed). (Terranova, 2004, p. 91)

Terranova suggests that this may no longer be the case. She argues that the fruits of this free, pleasurable labor of self-expression have been channeled within and structured by

capitalist business practices. The result is an increasingly blurry boundary between production and consumption, transforming individuals into “prosumers” and “producers” — people who create the content they (and others) then consume (Grinnell, 2009; Ritzer & Jurgenson, 2010; van Dijck & Nieborg, 2009; A. D. Williams & Tapscott, 2006). This new structure of public production results in a state wherein the pleasurable pursuit of self-expression and interpersonal interaction is also an opportunity for the generation of value for companies that have learned how to build software to glean commercially valuable information from those impulses. Terranova qualifies her argument by noting that free labor is not *necessarily* exploited labor; but the balance of valuation appears to increasingly skew in favor of corporations, rather than individuals.

I want to emphasize this uneasiness around the valuation of individual cultural production without conceptualizing individual activity as necessarily exploited. Certainly, the developers of gay-targeted social networking services make money from operating these apps by commoditizing user expressions and sharing this information with interested marketers; but we should not dismiss the unquantifiable (though decidedly extant) value derived by users in the course of engaging with these services.

A more concrete argument about the commercialization of online self-expression emerges in Alice Marwick’s research. In a networked media ecology increasingly structured by commercial interests, Marwick suggests that personal information — even when disclosed freely — has become a valuable commodity, and one which businesses are eager to channel into ever-more-lucrative forms (Marwick, 2005). In Marwick’s estimation, a central part of this process has been an increasing emphasis on unitary identities: online presences that are linked to offline markers of identity and stabilized

across time, space, and technical platforms. These unitary identities — predicted by Lawrence Lessig (1999)³ and referenced explicitly as a goal of Facebook by the service’s founder Mark Zuckerberg (quoted in van Dijck, 2013c) — become a way to compile increasingly detailed profiles of individuals with the intention of serving them personalized, targeted advertisements. The result, as Joseph Turow (2012) has described it, is the construction of a “daily you”: a networked view of the world fundamentally shaped by advertisers who are informed, in a historically unprecedented way, about the behaviors and preferences of individuals.

The result is a fundamental lack of control over the creation and circulation of individuals’ personal information. Far from the high degree of agency ascribed to individuals in Goffman’s account of dramaturgical identity performance, online identities are increasingly constructed *for* individuals, rather than by them, a phenomenon described by Mark Poster (1995) as “interpellation by database.” And, at its most extreme, this seeming lack of control has resulted in the emergence of reputation management services, which offer (for a fee) to carefully manage individuals’ identities online (Draper, 2014). Not only is expression itself monetized; but herein, the second-order product of expression (social reputation) becomes a target for commercial activity. Put another way, it’s telling that, 22 years after the initial publication of the famous *New Yorker* cartoon of a dog sitting in front of a computer with the caption, “On the Internet,

³ In this context, Lessig is concerned with digital identity as a sociopolitical construct tied into the enforcement of property rights, rather than identity as a facet of how individuals express themselves in networked social contexts. The theoretical takeaway is related: One paradigm for conceptualizing online identity posits that individuals only have one identity (be it a Facebook profile or a set of government records tied to a Social Security number); another suggests that identity is a multifaceted, contextually-situated assemblage of data. Lessig’s discussions of law on the internet privilege the former conception, while recognizing the practical complications posed by the latter. In my own work, and particularly in chapter 2, I focus almost exclusively on the latter.

nobody knows you're a dog," the magazine published its successor, in which one dog tells another, "Remember when, on the Internet, nobody knew who you were?"

The commercialization of the web at large is paralleled by the increasingly commercial logic of gay-targeted websites and services. Kate O'Riordan (2005) characterizes this shift as the transition from Usenet to Gaydar: from the wholly non-commercial space of Usenet newsgroups to the walled garden of a subscription gay portal. Where Usenet gave queer individuals an unstructured space within which to interact, sites like Gaydar channel gay sociability "into prescriptive identity menus" that are designed to serve the needs of service providers and marketers. Lisa Nakamura has characterized these interfaces as designed to enable "tribal" marketing strategies that address users as members of racial or sexual collectivities, rather than as individuals (Nakamura, 2002, pp. 122-123). And, particularly in the case of the most recent generation of gay-targeted services, the availability of increasingly detailed information about users (including geolocation data) enables service providers to construct highly granular pictures of app users as targets for marketing.

A second dimension of the commercialization of self-expression is an increasing tendency toward careful control and management of user behavior on the part of service providers. In "The Politics of Platforms," Tarleton Gillespie (2010) discusses how the term "platform" — used as a descriptor of online media of self-expression — obscures the structures of control that are built into the structure of mainstream services like Facebook and YouTube. Referring to a website as a platform, Gillespie suggests, leads us to believe that it is an "open, neutral, egalitarian and progressive support for activity." An ideal platform should be neutral to the content users display on it, and open to all types of

activity from different groups of people. In practice, however, the term “platform” has been applied to services that only rarely exhibit content neutrality and openness. In José van Dijck’s words, oftentimes a platform “shapes the performance of social acts instead of merely facilitating them” (van Dijck, 2013b, p. 29).

These structures of control and management are instantiated in terms of service, license agreements, and content management policies (discussed further below). While these policies are primarily intended to ensure a service’s compliance with national laws, or to provide indemnity for service providers in the event of illegal actions on the parts of users, it’s critical to recognize that they also have the effect of constraining what users are able to share on networked platforms. The emerging practice of content management (Gillespie, 2012) takes user expressions as a specific site for top-down control, applying often nebulous standards of decency and appropriateness to the wide range of content users share online. But, beyond an acknowledgement of their existence, these policies remain relatively unexplored in academic research, despite the fact that they structure at a fundamental level the everyday experiences of people online.

Online privacy and safety

Existing research into privacy and safety online has tended to follow one of three patterns: (1) Broad conceptual examinations of the new challenges to safety, privacy, and reputation created by new media; (2) Design-oriented studies of on-the-ground usability and accessibility in the implementation of privacy and safety features; and (3) Contestations of the status and utility of legal and policy documents pertaining to privacy and user rights.

A key component of many examinations of online safety is the question of control. Do individuals have agency in the use and distribution of their personal information, or do they frequently find it being viewed, shared, and acted upon without their consent or knowledge? A recurring theme in internet research is the suggestion that one of the consequences of the increasing popularity of digital/networked/social media has been a diminished capacity on the part of individuals to control their personal information and reputations. This diminished control is framed as a product of the essential attributes of networked media. danah boyd, for example, identifies four principal characteristics of networked media (what she calls “networked publics”) that differentiate them from the social technologies that preceded them: persistence, visibility, spreadability, and searchability (boyd, 2014a, p. 11). These affordances of media are well-documented, and in most cases obvious; but they make up the conceptual core of many of the anxieties about networked media documented by internet researchers. In practice, these affordances result in two outcomes that impact individuals’ ability to exercise control over their personal information: (1) Information tends to stick around after it’s shared; and (2) It’s fairly easy to share and re-share and re-re-share information after its initial appearance or disclosure in a networked setting.

Persistence, visibility, spreadability, and searchability have particularly pronounced consequences in the domain of individual reputation. In *The Future of Reputation*, Daniel Solove describes how digital technologies have transformed how individuals, groups, and societies negotiate identity and reputation.

As social reputation-shaping practices such as gossip and shaming migrate to the Internet, they are being transformed in significant ways. Information that was once scattered, forgettable, and localized is becoming permanent

and searchable. ... These transformations pose threats to people's control over their reputations and their ability to be who they want to be. (Solove, 2008, p. 4)

Our technologically-enabled capacity to remember, centralize, and share information has, in Solove's estimation, profoundly damaging effects on the autonomy of individuals, groups, and societies. In the face of persistent digital archives of even trivial embarrassing moments, we lose some of our ability to constructively author our own identities. The age-old practice of reputation management becomes significantly more challenging when gossip and criticism can become permanently attached to an individual's networked identity.

Solove's discussion focuses on identifying the correct balance of free speech and informational control. On one hand, we want control over our own reputation; but on the other, we want relatively unfettered access to the reputations of others — including the ability to translate personal wrongs against us into black marks on the public dossiers of others. This negotiation isn't specific to online interactions. However, when reputation emerged as a digitally networked practice — Solove historicizes this as taking place concurrently with the rise of blogging as a medium in the late 1990s — the scale of the negotiation changed. Reputation-affecting gossip spreads faster and wider online, and is harder to live down over time. The information architecture of reputation in networked contexts dramatically diminishes the ability of individuals to productively engage in self-protecting identity management. Permanent, virally-spread, legally unchallengeable “digital scarlet letters” are the new reputational norm — and, Solove suggests, if we commit the slightest social faux pas, any of us are at risk of being branded with one.

An important enabler of these shifts is an increased technological capacity to remember. Viktor Mayer-Schönberger chronicles networked information's inherent — and he argues unprecedented — tendency to persist after its creation.

Since the beginning of time, for us humans, forgetting has been the norm and remembering the exception. Because of digital technology and global networks, however, this balance has shifted. Today, with the help of widespread technology, forgetting has become the exception, and remembering the default. (Mayer-Schönberger, 2009, p. 2)

The precursors of this shift are technical: cheap digital storage and increased processing capacity make storage and recall inexpensive and easy. But, over time, the availability of technologies of easy remembering have trickled down into the social structure of software: interfaces subtly suggest that we “archive” information instead of deleting it, and users, in turn, begin to embrace this emergent norm of informational persistence. New tools spring up to do the work of curating these archives — from apps like Timehop, which surface our own actions on social networking services from years past, to ever-further-reaching search engines that index the increasing amount of information stored on networked systems (Halavais, 2009). Remembering, Mayer-Schönberger suggests, becomes a powerful default, rather than an active choice. He argues that our data management practices increasingly tend toward accessibility, durability, and comprehensiveness.

This overwhelming tendency to remember has clear upsides: We can more easily recall moments that brought us joy; we can counteract our tendencies to forget important details; businesses can operate more efficiently; knowledge is more readily available. But it also creates profound unease when technologies enable social or institutional memories of things we'd rather forget, or be better off forgetting. Forgetting — as individuals and

societies — enables progress; a technologically-diminished capacity to forget causes, in Mayer-Schönberger’s estimation, a trap of perpetual retrospection. By remembering too much, we become unable to move on from embarrassing, damaging, or regrettable moments. In essence, we lose the structures of forgetting that give individuals the ability to control their identities and reputations *despite* past embarrassments.

This lack of control becomes especially problematic when coupled with the visibility, spreadability, and searchability of networked media. Online, information about individuals is easy to access — not just by a person’s immediate community, but by (potentially) millions of individuals across the globe. The result is what Alexander Halavais describes as a fundamental shift in how identity works online:

The traditional view of identity in cyberspace is that it extends the trend of metropolitanism and cosmopolitanism: we simultaneously inhabit multiple selves and can easily step into alternative identities as we appear online. In contrast, search engines have thrown us back into village life in many ways. Public identities are often constructed out of what may be discovered via a search engine. (Halavais, 2009, p. 139)

Crucially, however, the “village life” Halavais describes can include multitudes of people. Individual reputations now unfold on a global scale.

Privacy is frequently positioned as a solution to the problem of unmanageable reputations; keeping personal information out of public view would seem to resolve the problematic tendencies highlighted by Mayer-Schönberger, Solove, and Halavais. A significant body of academic research takes up privacy as an entity that can be viewed relatively simplistically: certain technologies, interfaces, or services give individuals more privacy, while others take privacy away (boyd & Hargittai, 2010; G. Chen & Rahman, 2008; de Souza e Silva & Frith, 2010a; Madejski, Johnson, & Bellovin, 2011;

Stutzman & Kramer-Duffield, 2010). These studies examine particular individual behaviors — for example, the use of a “friends only” privacy designation for content shared on Facebook — and map out how certain choices result in either the broader or more narrow sharing of information within a social space. A frequent finding is that the developers of social networking services do not make privacy decisions easy enough to understand or access within a service’s interface, thereby resulting in less privacy for individual users. This design-oriented approach offers important insights into how interface designs can better reveal choices about disclosure and publicness to users.

Despite its utility, however, this approach has a number of important theoretical limitations by virtue of its emphasis on privacy as an outcome. While privacy may indeed be a motivating factor behind individual decisions, it alone does not offer a robust enough way to account for why individuals make particular choices about sharing and disclosure in different ways in different technological settings. Instead, we should focus on safety as an outcome, and position control and privacy as inputs that contribute to or detract from the ability of individuals and groups to be safe online. This perspective, which I take as a structuring element of my approach in this dissertation, has some grounding in existing research, although the term “safety” has not entered broad use in academic writing.⁴

One such reframing of privacy and control emerges in Helen Nissenbaum’s *Privacy in Context* (2010). The core of Nissenbaum’s argument is that privacy alone is not an analytically useful construct; instead, she offers the idea of “contextual integrity” as a more robust account of what actually makes individuals feel safe or unsafe online. Essentially, Nissenbaum suggests that “indignation, protest, discomfit, and resistance to

⁴ In contrast, “safety” is the dominant term used to describe these questions in industry settings.

technology-based information systems and practices ... invariably can be traced to breaches of context-relative informational norms” (Nissenbaum, 2010, p. 140); that is, she contends that where people admit to being anxious about their safety online, the issue often comes down to a feeling that the established norms of a given platform or online space have changed or been violated without the consent of the individuals in question. Her basic model posits that, in successful communication, the content of a given message (its “attributes”) is transmitted from one set of actors to another in accordance with a fixed set of accepted “transmission principles.” These transmission principles reflect the “context-relative informational norms” of the individuals or groups involved in a given communication. These norms could range anywhere from “this information cannot be used commercially” to “the subject of a transmission needs to give informed consent prior to the transmission taking place.” Crucially, context-relative informational norms are infinitely variable: Different spheres of social interaction have different transmission principles, reflecting the highly particular needs of individual actors. But the overall model remains the same: “Contextual integrity,” she writes, “is preserved when informational norms are respected and violated when informational norms are breached” (Nissenbaum, 2010, p. 140). While problematic in certain regards — namely, the permissiveness of “context” as a way to frame the sites of networked interpersonal interaction — Nissenbaum does the important work of framing privacy as a determinant of the overall goal of individual safety, rather than an end in itself.

A second, but equally significant, dimension of discussions of online safety is the contested status of policy documents such as privacy policies, terms of service, end user license agreements, and content guidelines. On a fundamental level, policy documents are

an attempt to articulate, in a legally-binding fashion, the terms of the relationship between users and service providers. These documents, often written using opaque, legalistic language (Gindin, 2009; Tasker & Pakcyk, 2008), have become a ubiquitous part of individuals' experiences of using software and the internet; but a number of significant questions about their function, enforcement, and utility remain unanswered. For instance, despite the fact that license agreements and terms of service have been in use regarding software for more than three decades, no new laws have been written to clarify their status in the United States. Instead, in a series of cases before federal courts, existing laws have been expanded to attempt to address the questions of notice, informed consent, and enforceability raised by software contracts (Dessent, 2002; Madison, 1998; Minassian, 1997; Pike, 2004).

Central among these questions is the issue of informed decision-making: Do users meaningfully understand the terms they're agreeing to? This issue has had significant resonance with regulatory bodies like the Federal Trade Commission (FTC). Recent research into regulatory behavior suggests a general level of skepticism towards the notion of a well-informed user, able to conscientiously enter into online contracts (Gindin, 2009). Users, this work argues, are not only uninformed, but also tend to eschew the practical steps they could take to become more informed. The FTC (2009) has acknowledged this tendency toward consumer inattention, independent of the length, clarity, or perceived importance of an agreement. It has, however, been unwilling to offer a firm normative injunction to users to actually invest the time and effort in understanding these agreements. The implication is that the agreements themselves, not users, are the problem.

Even if users took the 200 hours per year required to read every privacy policy, EULA, and TOS they encountered (Gomez, Pinnick, & Soltani, 2009), many would be unable to understand what they were reading. Despite attempts to educate consumers about online agreements, fewer than half of the respondents in a national survey indicated that the policies they encountered online are easy to understand (Turow, 2003; Turow & Draper, 2012; Turow, King, Hoofnagle, Bleakley, & Hennessy, 2009) — mostly because the policies are opaque, overly broad, and confusing (Anton et al., 2004; Gomez et al., 2009; Nissenbaum, 2010). These misunderstandings aren't helped by the fact that online terms of service are living documents, subject to frequent change with only minimal requirements for informing users about modifications. Even among individuals who *are* confident in their understanding of these policies, 66 percent have significant misconceptions about what service providers will or will not do on the basis of those agreements (Turow, 2003). The end result, as Turow points out, is that the resulting relationship between individuals, service providers, and marketers is not an equal one. The data gathering, storage, management, and dissemination practices of social networking services are often outside of the ability of individuals to audit or influence.

Finally, an emerging but vitally important branch of research into online safety has focused on physical safety as a counterpart to “virtual” questions of reputation management and privacy. Physical safety has been a recurring theme in public and academic discussion since the emergence of the internet as a popular medium, particularly with regards to the safety of children and adolescents (boyd, 2014a; Chun, 2006; Davidson & Martellozzo, 2013; Duncan, 2008; Turkle, 1995; 2011). The “Craigslist Killer,” a man who was accused of killing three sex workers who advertised

their services online, became a touchstone for anxieties about the commingling of online and offline interactions (LaRosa & Cramer, 2009). More recently, Grindr found itself at the heart of a series of controversies about the physical safety of its users, particularly in countries where homosexuality and same-sex sexual conduct is illegal (Mowlabocus, 2014a). I see this as an important recalibration of inquiry into the physical dimensions of online safety. In place of broad and often largely unfounded panics about hypothetical threats to the physical safety of individuals, we've begun to see a turn toward specific examinations of the design and data management practices of online services as they pertain to the protection of users. As I discuss below, the increasing popularity of locative and geosocial media has created compelling new possibilities for user expression and interaction; but these new technologies have and should continue to prompt important conversations about the responsibilities of developers and service providers for the welfare of their users.

The rise of locative media

A key element of the services in this study is the automatic collection and use of location information. Importantly, this is not an entirely new development. As early as the mid-1990s, a number of companies could provide, for a fee, information about an individual's postal code solely on the basis of their IP address. By the 2000s, cell phone tower triangulation provided location information about a mobile device's position that was accurate to within 100 feet. And, in the current generation of mobile phones and tablets, the inclusion of GPS receivers provides location data accurate to within 10-15 feet. Using this data, locative media connect users to relevant information and individuals based on

proximity. This allows users of mobile devices to access in real time information that is tailored to where they are in physical space.

In academic discussions of locative media, this has been characterized as a “hybridization” of physical and digital spaces — a union of offline geography with networked information (de Souza e Silva, 2006; de Souza e Silva & Frith, 2010b; Lee Humphreys, 2007; 2012; Sutko & de Souza e Silva, 2011). Importantly, these hybrid spaces prioritize neither their physical nor their digital dimensions; they constitute a new type of augmented reality (Manovich, 2006) in which online and offline contexts can seamlessly interact with each other. In the words of Adriana de Souza e Silva, who initially coined the term “hybrid space,” “flows of information that previously occurred mainly in cyberspace can now be perceived as flowing into and out of physical space, blurring the borders between both” (de Souza e Silva, 2006, p. 265).

In practice, this hybridity manifests itself in predominantly utilitarian ways. Many of the most popular location-aware applications are designed to navigate users through space or offer them relevant information about businesses in their proximity. Services like Google Maps, Urbanspoon, and Yelp use location information to allow users to make broad requests for information — for instance, a search for “Italian restaurants” — and receive immediately actionable results. Some services, like Foursquare (and its predecessor Dodgeball), couple this venue-based information with social components, allowing users to receive tips from friends and strangers about the locations they visit or receive notifications if their friends are nearby (Frith, 2013; Lee Humphreys, 2007; 2012). Importantly, these venue-based services — many of which require users to “check in” to locations they visit — anchor online interactions within particular physical

destinations. They add additional layers of information atop physical locations — for instance, recommendations or social data — but they do so without undermining the centrality of discrete venues.

Geosocial networking services — including the gay-targeted services I examine in this work — represent a radically different approach to the use of location information. While I examine the existing academic literature on gay-targeted geosocial networking services in the next section, I want to emphasize the novelty of the social model employed by these services. In place of privileging particular venues, geosocial networking services use location data as a way to connect users on the basis of proximity to each other, rather than proximity to a predetermined outside location. These apps take a range of forms, from apps targeted at finding friends or sexual partners (like Tinder and Grindr), to more open-ended services like Highlight, Ripple, and Secret, which allow users to participate, sometimes anonymously, in local conversations with other users. The shift from locative to geosocial media is a renewed emphasis on identifying and building social ties among users who are near each other, rather than solely on instrumentalizing the connection between people and the physical spaces they inhabit.

Gay

Handkerchiefs and subcultures

I want to briefly revisit the opening sentence of this dissertation: the contention that gay social media has moved from the margins into the popular spotlight. In the overall context of Western popular culture, where an increasing number of massively popular television shows (like *Glee* and *Modern Family*) feature gay, lesbian, bisexual, and (to a

lesser degree) transgender characters, this seems like a fairly unproblematic claim to make. The general tendency toward the mainstreaming of gay culture in the media has been productively traced by, among many others, Katherine Sender (2003; 2004) and Larry Gross (2013). These works suggest that, through mass media texts, LGBT individuals have not only themselves become more visible, but the “codes” of gay culture — the rhetoric and style of gay characters — also become more accessible to mainstream audiences.

These encoded practices of subcultural performance take a wide range of forms, and have important histories for gay men in the West. An especially significant pre-digital example is the “hanky code,” a system of visual identification of sexual preferences using different configurations of colored handkerchiefs. In 1970, a journalist in the *Village Voice* joked that gay men could resolve the potential ambiguities of determining whether someone was top or bottom (that is, the penetrating or penetrated partner in anal sex) by wearing colored handkerchiefs to broadcast their desires. A hanky in the left pocket indicates that its wearer is a top; a hanky in the right, bottom. And, subsequently, the visual language of the hanky code blossomed into a robust vocabulary (Levine, 1998, p. 66; Patton, 1990, p. 46). The hanky code engaged with the precarious state of post-Stonewall gay visibility by putting sexuality in plain sight: the intricacies of an individual’s sexual predilections were publicly available to any passer-by, so long as that passer-by was fluent in the variegated parlance of back-pocket handkerchiefs. The plausible deniability offered by the hanky — it remained, in the end, a handkerchief — balanced the desires of gay men to form networks of sexual contacts with the need to keep their (then still illegal) sexual practices out of the straight public’s eye.

Of course, these practices of encoded expression are not unique to gay male subcultures. Dick Hebdige's influential work on British youth subcultures emphasizes the significance of particular performative styles — for instance, clothing, tastes in music, manners of speaking and walking, the use of particular recreational drugs, and so on — for the expression of group identities (Hebdige, 1979). Hebdige draws on a wide range of theoretical influences, including Raymond Williams (1961) and Stuart Hall (1973), in arguing that the mundane elements of self-expression he observed with British youth were in fact important parts of a system of meaning-making tailored to the particular experiences and needs of those groups.

A key affordance of subcultural style, both in Hebdige's analysis and more generally, is the ability to quickly and easily express information about group identification and identity. In Hebdige's work, wearing a particular style of jacket or torn jeans could immediately make an individual legible as a member of a particular subcultural group. For gay men in the 1970s and 80s, the hanky code delimited specific sexual practices by expressing them symbolically with colored handkerchiefs. And, over time, these codes evolved within gay male culture to enable individuals to share yet more detailed information about their bodies and interests. For instance, the somewhat tongue-in-cheek "Natural Bears Classification System" gives individuals the ability to share properties like the length and bushiness of one's beard (from B0 to B9), height, weight, and body hair through the use of letters, numbers, and symbols (Donahue & Stoner, 1997), presumably with the intention of making those properties more easily expressible in online spaces like chat rooms or message boards (Campbell, 2004; Wright, 1997; Wright & Wehrle, 2001). In each case, these forms of encoded expression allow for easy

signification of identity through the overt display of an agreed-upon symbol. I see these coded expressions of identity as a crucial theoretical predecessor to subsequent forms of gay male identification and identity work — including the creation of online profiles I examine more specifically in chapters 2 and 3.

More recently, however, the use of particular stylistic signifiers or subcultural identifications within American gay culture has been met with a turn toward authentication. Authenticity discourses suggest that not just any individual can legitimately claim to be a part of a particular subcultural community merely by adopting its stylistic or rhetorical signifiers. This tendency reveals itself clearly in a recent attempt by one individual to quantify gay subcultures through the use of an online survey. The Gay Cliques Census (Hafertepen, n.d.), as its author termed it, relies on the disclosure of quantifiable embodied attributes to ascribe subcultural identities (such as “bear,” “otter,” or “twink”⁵) to individuals. In this case, we see not the openness of subcultural identification through style that Hebdige discussed, but rather an attempt to authenticate bodies as belonging to particular communities through an assessment of embodied characteristics.

This combination of authenticity and positivistic approaches to the body creates a particular problem for the identities and bodies that online gay communities historically have embraced. We can begin to reconcile these theoretical tensions — between the historical openness of gay identification on one hand, and the tightly constrained

⁵ These terms are used as broad labels for certain types of bodies. Bears tend to be larger-bodied and hirsute. Otters are hirsute but slim. Twinks are generally young and have little or no body hair. The particular dimensions of each term — and the boundaries between terms — have traditionally been left relatively undefined.

establishment of identity categories online on the other — by turning to Lionel Trilling’s discussion of sincere self-presentation. Proper sincerity, Trilling writes, requires that “we play the role of being ourselves” (Trilling, 1971, p. 11) — that who we say we are, as a question of social construction and interpersonal engagement, is judged against what we are, as a question of embodiment. As John Jackson puts it,

Authenticity attempts to domesticate sincerity, rein it in, control its excesses. It demands hard, fast, and absolute sure-footedness, whereas ... sincerity wallows in unfalsifiability, ephemerality, partiality, and social vulnerability. Sincerity highlights the ever-fleeting “liveness” of everyday ... performance that cannot be completely captured by authenticating mediations of any kind. Where authenticity lauds content, sincerity privileges intent. (Jackson, 2005, pp. 17-18)

The search for certainty of “correct” identity results in holding one’s sincere performance of self to some objective standards of authentic identity. We seek in notions of authenticity some way to disentangle the essential problem of sincerity: individuals assuming identities that may or may not be real, and that accordingly threaten the very structure of the groups they identify with. Decades of contestation have complicated notions of what participation in gay communities looks like, making the boundaries between subcultures fundamentally ambiguous. Sincerity is the enabler of that ambiguity; authenticity is a corrective.

Gay communities, both online and offline, are simultaneously brought together and held apart by these two concepts. The semiotic openness of online self-expression speaks to a form of sexual sincerity that looks beyond bodies and towards a common search for community. But the fact that gay identity originates in a feeling of difference, of separation from mainstream culture precisely because of embodied characteristics of self, implies at least some relationship to ideas of bodily authenticity. The act of creating

a profile on a social networking service asks users to reflect on their sincere performance of self as grounded in the authentic “stats” of their body, as well as socially-negotiated categories.

Getting off online

Gay men’s use of networked media significantly predates the emergence of mobile applications. In the 1980s, the pre-World Wide Web French telecommunication platform Minitel played host to a variety of gay and lesbian services, including discussion boards, gay-friendly business listings, and erotic personals messages known as *messengeries roses* (Duyves, 1993; Livia, 2002). Around the same time, the Usenet group soc.motss (where “motss” stands for “members of the same sex”) became an important site of what researchers have termed “cyberqueer” community in English-speaking countries (O’Riordan, 2005; 2007; O’Riordan & Philips, 2007). As Nina Wakeford (2002) notes, these queer communities achieved success despite — and perhaps because of — pervasive homophobia in many early networked communities; where LGBT individuals encountered opposition on mainstream fora (like the popular net.singles newsgroup), they formed separate, distinctively queer online spaces.

By the 1990s, Internet Relay Chat (IRC), as well as commercially-supported chat rooms on services like Microsoft’s MSN and America Online, became popular venues for gay men to meet and interact on the internet. In 1994, for example, three of the top 10 most popular IRC chat rooms (according to a list compiled by *Wired*) were targeted at gay men: #men4men, #MenWhoWant2MeetMen, and #YoungMen4Men (Wakeford, 2002, p. 119). The medium of text-based chat offered gay men a chance to not just

interact with each other online, but to form niche communities based on their specific interests and embodied characteristics. John Edward Campbell's discussion of gay men's use of IRC is instructive: Campbell suggests that participants in different chat rooms emphasize different elements of their bodies or interests depending on the stated context for that particular chat room (Campbell, 2004). For example, Campbell found that the participants in the #gaymuscle chat room (a community based around an interest in the muscular gay body) emphasized certain parts of their bodies, such as the circumference of muscles and their percentage body fat. The chat rooms Campbell studies are spaces where users have the opportunity to explore different subcultural communities without necessarily embracing those communities offline. The multiplicity not only of identity categories, but also available social contexts on the internet allows for a flexibility of individual expression that, Campbell suggests, is often foreclosed by mainstream discourses of sexuality. The anonymity of the text-based IRC protocol gives users the ability to present a more specifically-constructed version of themselves, better suited to the complex interplay of motivations, interests, and bodies brought together under the banner of "gay identity."

As broadband internet connections became more widespread in the early 2000s, a new generation of higher-bandwidth gay-targeted web portals and social services (such as Gaydar, Manhunt, and Gay.com) rose to popularity, offering users graphical profiles and the ability to easily exchange photos (Cassidy, 2013; Gosine, 2007; Light, Fletcher, & Adam, 2008; Mowlabocus, 2010a; 2010b). In many cases, these early gay networking services were characterized (in both academic and popular accounts) as sites of online

cruising.⁶ From bathhouses to erotic movie theaters to public parks and restrooms, academic scholarship on gay male culture has described the importance of cruising as a sexual and social activity (Capino, 2005; Colter, Hoffman, Pendleton, Redick, & Serlin, 1996; Delany, 2001; Laud Humphreys, 1975; Maynard, 1994). It's unsurprising, therefore, that studies of gay men's use of networked media have drawn parallels between gay-targeted online services and offline acts of cruising. As David Gudelunas puts it,

The act of cruising has moved online and to mobile phones, but the effect is the same. Gay men can still hail one another within anonymous crowds in order to both solidify their real and imagined social networks as well as find partners for practical, sexual pleasures. (Gudelunas, 2012b, p. 14)

Gudelunas, as well as Mowlabocus and many of the authors of the studies of Grindr I discuss below, suggest that, whatever the technological novelty of these apps, they are little more than online cruising areas. A reliance on cruising as a metaphor for gay social networking reduces the complexities of these services to a simple recreation of analog forms of interaction. This both neglects some of the crucial characteristics of gay-targeted social networking services, and dismisses non-sexual uses of gay-targeted networked platforms.

Importantly, the second generation of gay-targeted social networking services — platforms like Manhunt and Gaydar — reinforced the centrality of images in gay men's online presentations of self. In the early 2000s, the emergence of cheap, easy-to-use camera peripherals (webcams) gave rise to a host of services built to allow users to easily share video clips or stream live content to each other — a practice referred to in academic

⁶ "Cruising" typically refers to wide range of offline techniques for finding sexual partners, typically for anonymous or semi-anonymous one-time sexual encounters.

research as “televideo cybersex” and popularly dubbed “camming” (Döring, 2009; Shaw, 1997; Waskul, 2002). Waskul (2002) begins his study of cybersex with the assertion that it, unlike text-based interactions like IRC, camming constitutes a more directly embodied experience between individuals. Embodiment, he suggests takes place at the moment where one participant sees the body of another — a claim that stops short of asserting the separateness of online and offline presences, but nevertheless establishes a hierarchy of online presence by the “directness” (that is to say, visuality) of a particular interaction. Sharif Mowlabocus (2010b) makes a similar claim about the role of photography in online dating sites. The experience of seeing and being seen — rather than reading and being read — is understood in these accounts as more directly or satisfyingly erotic for the individuals involved.

Mowlabocus describes this form of digital eroticism as a type of “cybercarnality” that bridges virtuality and embodiment in online spaces (Mowlabocus, 2010a).

Mowlabocus suggests that embodiment in gay male digital spaces takes place in large part through a process of pornographic remediation, and insists that gay embodiment in digital spaces (as opposed to other types of social networking) is “intimately tied up with the structures of looking and of consumption” that are most characteristic of how people consume pornography (2010a, p. 81). A similar, but broader, claim is made by Shaw (1997), who notes that the “extraordinary visual bias” of gay culture results in an overwhelming prevalence of images and video in gay communities online. This visual bias, both in gay communities and more generally, has prompted a recurring set of questions about authenticity in online self-expression. Many studies of social media — and especially of services targeted at forming romantic or sexual relationships — have

focused on the veracity of photos and profile text in self-presentation (Ellison, Hancock, & Toma, 2012; Ellison, Heino, & Gibbs, 2006; Hancock & Toma, 2009; Toma, Hancock, & Ellison, 2008). In these studies, photos operate as what Mowlabocus (2010b) calls a “passport” into the world of online relationships: meaningful participation requires some form of visual authentication of self.

I take this emphasis on the visual as a point of departure in my analysis. While images are a central part of how profiles are constructed on Grindr, they constitute one datum among many others in a user’s networked presence. Perhaps, as Mowlabocus, Shaw, and Waskul suggest, images are indeed the most important part of this presence — a claim that is at least in part validated by the centrality of images in the interface design of Grindr. Nevertheless, this study examines the ways in which images are augmented and anchored by a wide range of other pieces of personal information, from quantitative “stats” about a user’s body to categorical identifications with terms like “bear” or “twink” to automatically-collected data about a user’s geographic location. Put more generally, we might restate Shaw’s claim about the “visual bias” of gay culture as an emergent data bias: a fixation on information collection and display as part of the experience of online gay sociability.

Gay-targeted geosocial networking

Before turning specifically to the small but quickly-growing body of academic research focusing specifically on gay-targeted geosocial networking services, it’s worth acknowledging that geosocial networking applications were not the first use of locative and proximity-based technologies for the purposes of interactions among gay men. In his

discussion of digital cruising practices, of which Grindr is one example, Mowlabocus makes reference to the technique of “Bluejacking,” a term used to describe the practice of exchanging short, often explicit text messages and photos over Bluetooth connections (Mowlabocus, 2010a, p. 191). This use of Bluetooth, a short-range wireless communication protocol integrated into many mobile devices including pre-smartphone mobile phones, employed the specific affordances of Bluetooth (namely, its low range) to facilitate only connections amongst those users in immediate proximity of each other.

(The canonic example of Bluejacking in Mowlabocus’s work is on a subway train.)

While Bluejacking never attained significant popularity — one imagines this is due to the difficulty associated with finding and establishing these short-range connections, along with the inherent security risks of allowing one’s phone to connect over Bluetooth with any nearby device — it speaks to a tradition in gay communities of innovatively employing new technologies to facilitate interpersonal interactions. And, more specifically, Bluejacking — like other spatially-situated techniques of cruising — emphasizes the importance of location in gay male self-expression and interaction.

The emerging body of research focusing on geosocial networking applications has continued this focus on the interplay between location and self-expression. Roderick Crooks (2013), for example, discusses the role of Grindr within the particular geographic space of West Hollywood. Grindr users, including Crooks himself, engage with the app as a way to not only meet individuals in their immediate vicinity, but also solicit location-specific information (such as news and restaurant recommendations). David Gudelunas (2012b) likewise stresses the importance of location, noting that the use of geosocial apps while traveling was a recurring theme amongst his informants. Location also informs the

ways in which users construct their profiles: Blackwell et al (2014), for example, suggest that while locality affords users of geosocial apps the possibility to easily interact face-to-face, it also prompts significant tensions and anxieties around strategies of self-presentation when the face-to-face encounter is so readily at hand. Birnholtz et al (2014) make a related claim, suggesting that patterns of self-presentation correspond with particular localities; individuals on a college campus choose to construct their profiles differently than individuals in a crowded urban area.

Another prominent theme in existing research on Grindr is the limited amount of information users are able to express in a profile. Approaching the app from the perspective of human-computer interaction, Birnholtz et al (2014) discuss the ways in which users manage “potentially stigmatized identities” (namely, seeking casual sex) through the selective disclosure of specific types of information. While the ability for users to share detailed information about their sexual interests is constrained by the content management policies of services like Grindr, Birnholtz et al nevertheless identified recurring phrases (such as “DL” and “discreet”) that, in their estimation, corresponded with users’ desires to use the app for particular behaviors.

A significant number of academic studies of these services focus on health behavior, following one of two dominant patterns. One cohort of studies take Grindr use as an independent variable, examining the correlation between an individual’s use of Grindr for sex-seeking and a wide range of sex-related health behaviors, such as condom use and regular HIV testing (Lehmiller & Iorger, 2014; Rendina, Jimenez, Grov, Ventuneac, & Parsons, 2013; Rice, Holloway, & Winetrobe, 2012; Winetrobe, Rice, Bauermeister, Petering, & Holloway, 2014). Others examine the usefulness of networked

platforms like Grindr for spreading health-related messages within a particular community (Rosser et al., 2011) — a practice whose effectiveness is still contested but which has been adopted by a number of urban public health departments. In each case, these studies suggest that Grindr use constitutes a meaningful evolution in the sexual and health behaviors of men who have sex with men — that men who are on Grindr are consistently, quantifiably different from men who aren't on Grindr with regards to their attitudes towards relationships, sexual practices, or health behavior. While the actual results of the studies are more ambiguous — most are unable to establish a strong link between Grindr use and less safe sexual behavior — they nevertheless assert that Grindr users are a “high risk” group that ought to be targeted by public health initiatives.

Finally, it's important to recognize that not all gay-targeted social networking services are equivalent. Gudelunas (2012b) suggests that, amongst his informants, some maintained profiles across as many as 7 to 10 different social networking services, using the particular affordances of different platforms as a way to express different elements of their bodies, identities, and interests. Specific services, such as Scruff, may target themselves more narrowly to certain subcultural communities, such as bears, while others take a one-size-fits-all approach (Roth, 2014). Users actively negotiate these differences, selecting services that best align with their particular needs.

More generally, this plurality of competing services and user motivations should remind us that we cannot meaningfully make general assertions about what apps like Grindr are “for.” While a number of scholars have branded these services “hookup apps,” prioritizing their use for seeking casual sex, I want to stress that sex, dating, and relationships are three among a potentially limitless set of possible use cases for gay-

targeted geosocial media. These services can and do offer opportunities to gay men to find friendship with other gay men. They can be a source for information about local gay venues and events for users visiting a new city or country. They already have been employed by public health groups as ways to spread information about safer sex, and by political advocacy organizations as ways to get out the vote or share information about ballot initiatives. Grindr even includes professional networking as an option in a menu of choices for what users are “looking for” on the service. My goal, in focusing on the structure, design, and management of these services, is to identify the potentialities of expression and use they enable, rather than ascribe a fixed set of uses to them.

Method

This dissertation adopts an intentionally eclectic and polyvalent approach to collecting, analyzing, and presenting information about its objects of study. Across each of its chapters, I introduce evidence drawn from hands-on experiences with applications, interviews with users and software developers, blog posts, software support pages, app store reviews, marketing materials, developer guidelines, technical documentation, and mainstream press articles. My discussion of Grindr, and of gay-targeted social networking more generally, moves between analytic registers focused on micro-level user experience and macro-level economic and technical structures — all of which, I argue, are essential components of the broad sociotechnical paradigm that constitutes gay-targeted networked media. Necessarily, this requires a set of methodological approaches suited to organizing and systematically evaluating multiple forms of data. In order to position this study within existing methodological and empirical traditions, this section

(1) identifies the two dominant existing approaches to studying software, which I term the “micro” and “macro” traditions; and (2) identifies the elaborations on existing traditions offered by the specific methodological approaches used in this dissertation.

At a schematic level, we can differentiate between micro and macro approaches to studying software by the emphasis they place on particular users. The micro tradition is focused on establishing how particular platforms or applications are used by particular individuals or groups. These analyses begin with evidence drawn from individual users, in the form of interviews, walkthroughs, laboratory experiments, or usability surveys. When successful, these accounts are necessarily socially, culturally, and temporally specific; they articulate the contours of a particular moment in the operation of a particular platform, for a particular set of individuals. The macro tradition, by contrast, begins with data that is not individually-specific, in the hopes of identifying the properties of sociotechnical systems that might impact the on-the-ground experiences of multiple different constituencies of users. While these studies tend not to rely heavily on empirical data, they draw on textual and visual analysis to inform the schematic, often speculative accounts of software they offer. When successful, these views of software offer generalizable explanations of how the recurring structures of sociotechnical systems enable or constrain broad patterns of user behavior.

Recent work by Mirko Tobias Schäfer (2011) clearly articulates the conceptual differences between these approaches. Schäfer draws a distinction between two types of users that are “present” in any piece of software: what he terms explicit and implicit users. An explicit user is the “real” user, interacting with a particular application. Ethnomethodologies that observe users in situ and interview-based approaches offer

access to this type of user. In my typology, micro approaches tend to be focused on explicit users. An implicit user, by contrast, is the user imagined by an application's developers: the person and types of interactions they have in mind when designing the interface features of a service. But, rather than ask developers directly who their implied users are, many of these studies adopt a methodological approach aligned with the tradition of critical code studies, introduced in works like Galloway's *Protocol* (2004), Manovich's *Software Takes Command* (2013), and Fuller's *Behind the Blip* (2003). An analytic focus on implicit users aims to highlight the structures and constraints created and reinforced by technical systems as they affect a wide range of differently-motivated, experienced, and engaged users. In my typology, this would be characteristic of a macro approach. The central contention of the macro approach is that individual users are fickle; software itself is less so.

Arguably, the macro tradition is the branch of software studies that has been most comprehensively — or at least, exhaustively — discussed on the order of abstract methodological considerations. Matthew Fuller (2003), for example, suggests that our understandings of software must go “behind the blip,” setting aside the functionalist logic of user experience studies in order to allow software to stand in for itself. (He stops short of articulating specific sources of data that enable us to do so.) Alexander Galloway (2004), in turn, offers an archaeology of the connective protocols underlying networked media, insisting on “the network” as both a material entity and a metonym for social processes taking place online. Geoff Cox (2013) urges us to understand programming as a performative utterance — that code says what a technical system will do, then actually does it, an echo of Lawrence Lessig's famously pithy “Code is law” (Lessig, 1999). Lev

Manovich (2013) writes of “cultural software” — technological artifacts that meaningfully carry within them the “atoms” of their users’ cultures. Adrian Mackenzie (2006), citing George Marcus (1995), takes software as a multi-sited assemblage of code itself, users, technical systems, and social contexts. Each of these works articulates a vision of software that goes beyond “mere” user experience, and aims for a more durable account of the underlying substance of software: what many authors refer to as “code,” but which goes far beyond just the languages used to program a particular application. Herein, “code” seems to be the analytic anchor for accounts of software that often focus on what a protocol, program, or platform *means*, rather than what it *does*.

The micro tradition, by contrast, is determinedly functionalist. These approaches have tended to focus narrowly on what software does, and how particular, empirical facets of software and user behavior come together to enable that activity. Methodologically, this entails an exacting account of either/both an application’s interface or/and the motivations and behaviors of existing users. The methodological framework of software “walkthroughs” has become a common way to conceptualize how to gather and organize micro-level information about software and user behavior. Walkthroughs, notes Mia Consalvo (2003), are commonly employed by video gamers as a way to share information about the narrative and ludology of a game, as part of an intertextual framework of participation and consumption within gamer communities. Walkthrough videos chronicle both the “fixed” structure of a game, as well as the more fluid activities of gamers as they engage with the game software. The resulting text is an encapsulation of both the affordances of the game, and the behavior of users with regards to those affordances. In game studies, and in software studies more generally, the

construction of walkthroughs has become a common way to systematically outline the operation of novel sociotechnical systems. A related approach, outlined by Boland (2002), Cole and Avison (2007), and Burnett et al (2013), offers a “hermeneutics” of software on the basis of a similarly exacting walkthrough. In these cases, researchers position themselves as supposedly naive users, and document in exacting detail their interactions with software, in the hopes of arriving at a comprehensive and “objective” account of the micro-level operation of software.

A different version of the micro tradition of software studies relies on interviews and surveys with users themselves to contextualize findings about software interfaces and affordances introduced by the authors of a study. This is the dominant approach within the existing body of scholarship on gay-targeted social networks. Many studies of these services focus on establishing the uses and gratifications of a given app or platform for its users. These studies primarily employ one-on-one interviews, focus groups, and surveys to describe, both qualitatively and quantitatively, why gay men choose to use particular apps or websites to interact with each other (Blackwell et al., 2014; Crooks, 2013; Gudelunas, 2012a; 2012b; Raj, 2011; Van De Wiele & Tong, 2014). While valuable, these studies treat the technosocial configurations of gay-targeted social networking services as relatively stable entities whose dimensions can be understood simply by asking users what they think of them, or why they made particular choices. Revealing though these descriptions of user behavior may be, they often substitute the accounts of particular users for a more fundamental examination of the temporally and contextually situated properties of the technosocial systems that inform them.

Some recent approaches have begun to hybridize the micro and macro traditions, examining software as both an applied practice (in an individual and institutional sense) as well as an ideal type. For example, José van Dijck’s method of platform analysis (van Dijck, 2013a; 2013b; van Dijck & Poell, 2013) uses empirical data to establish an ideal type for the sociotechnical systems being studied. Across each of her case studies (bounded platforms, such as Facebook, Twitter, or Flickr), van Dijck uses granular information about the history, code, financial structure, and user experience of a given platform to identify the broader principles of its operation. van Dijck’s approach offers two critical elaborations on both the micro and macro traditions of software studies. First, van Dijck identifies sources of “micro” data that were not captured by existing methodological approaches such as walkthroughs. By incorporating financial and political considerations into micro-level analysis, van Dijck recognizes the ways in which concrete, empirical factors exogenous to a piece of software can nevertheless have a direct impact on its operation. Second, van Dijck positions micro-level data as a source of macro-level insight into the identity and impact of the platforms she considers. This approach allows van Dijck to ground broad analyses of networked media in the particularities of individual platforms — a hybrid analytic approach that yields compelling insights while eschewing the largely unempirical poetics typically associated with macro-level software studies. While valuable, van Dijck’s framework largely offers pointers towards the types of questions a researcher should ask, rather than a concrete articulation of the tools a researcher should use to answer them.

The approaches outlined in this section are my answer to the open question of research methods offered by van Dijck. In selecting the appropriate methods for this

study, I want to emphasize what I see as need to move beyond descriptive, micro-level research on one hand, and over-broad, largely un-empirical assessments of software on the other. Rather than examining what the uses of a given platform are for particular individuals (or groups of individuals), I want to reveal how particular applications and services enable or constrain the possible uses that might be undertaken by a diverse body of individuals. Instead of focusing on how a particular service has worked in the particular praxis of a group of individuals (analogous to traditional audience reception studies in communication research), I want to identify the design principles, values, and ideologies that emerge through the temporarily stable entities that are social networking apps and websites. To do so in in this project, I use a combination of ethnographic participant observation, textual analysis, interviews, and what I term “interface analysis.”

Interface analysis

A principal dimension of this study’s methodological approach involves the close analysis of software interfaces. While a number of theoretical and empirical texts have done the critical work of mapping the unfamiliar sociotechnical terrain of identity work on digital and mobile media (boyd, 2014a; Cassidy, 2013; de Souza e Silva, 2006; de Souza e Silva & Frith, 2012; Duguay, 2014; Farman, 2012; Gillespie, 2010; Lee Humphreys, 2007; 2012; Mowlabocus, 2010a; Papacharissi, 2002a; 2002b; R. Schwartz & Halegoua, 2014; van Dijck, 2013b), few have offered a granular account of software interfaces themselves. Much of the existing research into online platforms seeks out the commonalities between disparate services (Papacharissi, 2009; van Dijck, 2013c), drawing out the “essential” elements of sociotechnical system, without regard for

the specific affordances, constraints, and practices associated with individual interfaces. These schematic accounts are helpful, but are only a starting point. We have developed macro-level theories of digital media without first doing the necessary micro-studies of software. In part, my methodological approach is one of disaggregation (Langlois, McKelvey, Elmer, & Werbin, 2009): of separating applications and platforms into their constitutive parts — infrastructures, interfaces, and interactions — in order to more specifically parse the work of digital interfaces.

The traditional definition of “interface” stresses contact between two different objects or systems: a point where discrete entities interact with each other through a mutual medium. At the most schematic level, therefore, we can describe interfaces as what Peter Galison (1997) terms a *trading zone*: a point of contact between systems that rely on shared access to certain symbolic resources to enable productive interaction. The visual and symbolic logic of interfaces — icons, hyperlinks, drop-down menus, and so on — provide the resources necessary for humans (users) and computers (applications) to productively interact. But, as Galison points out, the fact that a trading zone is functional does not mean that the interactions it enables are wholly agreed upon by both parties; an interface may constrain what users are able to say or do by failing to offer the right menu of options; and users, in turn, may put interfaces to use in ways that their designers neither anticipate nor authorize. To borrow from Stuart Hall’s (1973) classic description of media, interfaces both give users the ability to encode information in ways that make it legible to computer systems, and to decode data in ways that render machine knowledge useful to human audiences. As Hall acknowledges, all acts of encoding and decoding are lossy; perfect translation between audiences (in this case, users and computers) is all but

impossible. In this sense, we can conceptualize interfaces as what Bruno Latour (2005) calls a *mediator*: a highly particularized entity that transforms, distorts, and modifies the information it carries. The task of interface analysis, therefore, is to reveal the mechanisms by which these processes of translation, transformation, and distortion take place.

To gain access to the symbolic logics of software interfaces, I treat interfaces as texts to be closely read. This approach has been commonly deployed in research identifying with the traditions of software studies (Berry, 2011; Fuller, 2003; 2008; van Dijk, 2013b) and critical code (Chun, 2006; Cox, 2013; Galloway, 2004; Hayles, 2004; Schäfer, 2011), but has rarely been systematically described as an analytic technique. In an effort to make interface analysis more concrete as a method, I divide my approach to close reading into two tasks. First, what I term a *syntactic* approach to interfaces examines the details of how a particular interface is constructed. Syntactic questions might include:

- In what order are interface elements (e.g. text boxes, images, buttons, etc) presented to users?
- Which items are visible or hidden?
- Which interface features are highlighted to users?
- To what extent are design elements and interface features explained to users?
- What are the design principles (e.g. aesthetic choices, usability concerns, etc) that factor into a particular interface?
- What are the default states, or automatically assigned settings, of interface elements?

These syntactic questions correspond with what James Gibson (1986) described as objects' "affordances": that is, the ways in which the material properties of an object influence how actors subjectively perceive the uses of that object. Gibson's perspective, emerging out of psychology and ecology, closely parallels work in the history of science that asserts that artifacts in themselves have politics (Winner, 1999). The rhetoric of affordances, which I use throughout this work, is intended to indicate a mutuality between and mutability to both objects and individual intentions (Majchrzak, Faraj, Kane, & Azad, 2013). The syntactic properties of interfaces — that is, those architectural properties of a system that are fixed from a user's perspective — frame the possibilities users have for acting within those interfaces (Hutchby, 2001); but these fixed properties do not foreclose the existence of multiple possibilities for individual action (Cirucci, 2015; McVeigh-Schultz & Baym, 2015; Nagy & Neff, 2015; Oudshoorn & Pinch, 2003; Schrock, 2015). This perspective recognizes the ways in which technosocial systems offer relatively fixed, finite capabilities to users, even as it rejects the deterministic suggestion that users only act in the ways suggested by those capabilities. Put more directly: Examining the syntax of an interface reveals that interface's affordances for user action.

Second, a semantic approach to interfaces investigates how the fixed characteristics or affordances of an interface come together to construct a particular subjective user experience. Semantic questions might include:

- How does the design of an interface reflect the intended "identity" or "personality" of an application?

- How does the presence or absence of certain interface features enable or constrain a particular form of user behavior?
- What are the underlying ideologies (Chun, 2006; Nakamura, 2001; 2002), values (Flanagan, Howe, & Nissenbaum, 2008; Friedman & Nissenbaum, 1996; Nissenbaum, 2001), or norms (Gillespie, 2010; van Dijck, 2013b) that are reflected in the design of an interface?

Using these questions as a guide, I examine both static images of interfaces (screenshots) as well as the live usage of applications and services. I draw these screenshots from the marketing materials and press kits of each of the services I examine, as well as from my own usage of these apps. Where interface designs have changed substantially over time, I have maintained an archive of screenshots designed to highlight the relevant changes.

This approach to studying software emphasizes the characteristics of a given application or service that are common to all users, rather than solely those features highlighted by particular users in interviews. While all analyses of software, including my own, are subject to interpretive biases, my goal in beginning with a set of specific, descriptive questions is to compile a meaningful account of a given interface before turning to the evaluative task of understanding how the elements of an interface construct particular user experiences.

Participant observation

Static images and screenshots only reveal a limited slice of the overall experience of using an application or website. A wide range of dynamic or less tangible characteristics — the responsiveness of the interface; how software reacts when the user taps a button;

unexpected, intermittent patterns of user behavior; conventions that emerge only in practice; constraints that only reveal themselves when a user encounters them during everyday use — can only be identified through participant observation. Accordingly, I have also employed ethnographic participant observation as a data collection strategy. Between 2009 and 2013, I was an active user of each of the services I study, engaging with each almost daily in a wide range of geographic settings (from the United States and Western Europe to the Middle East).

Even as this study is not, at its core, an ethnography of gay-targeted social networking services, I want to briefly highlight some of the methodological considerations and complications associated with using ethnomethodology as a way to study software and the internet.

In part, my approach draws on the emerging body of research into virtual ethnography (Boellstorff, 2008; Boellstorff, Nardi, Pearce, & Taylor, 2012; Hine, 2000) as a way to guide my approach to data collection and analysis in online environments. Superficially, these studies seem like an uneasy fit: There is undoubtedly something exotic — by design — about the virtual environments of World of Warcraft, EverQuest, and Second Life; we see them, in the work of Boellstorff, Nardi, and Taylor, rendered in all their alien splendor. But embedded beneath the fantastical exteriors of these worlds is a computer application not entirely dissimilar from the word processors, e-mail clients, and web browsers we interact with every day. Virtual worlds are, at their core, software; and the methodological techniques used in virtual world ethnography — for instance, detailed field notes, systematic exploration, informant interviews, and experiential immersion — can just as readily describe an ethnographic inquiry into any other

networked environment, including, critically, a gay-targeted social networking application. I adopted many of these conventions (which, arguably, could just as readily describe an offline ethnography as a “virtual” one) in documenting and analyzing my own use of these services, and they frequently serve as an empirical anchor for the interface and textual analysis I undertake throughout this dissertation.

Despite the fact that these ethnomethodological approaches are, on the whole, quite traditional, framing my study of Grindr as akin to a virtual ethnography raises a number of substantial methodological and ethical questions. Principally, even as I treat the Grindr application as a relatively stable “virtual” entity, it’s crucial to recognize that Grindr does not operate independently from physical, decidedly un-virtual geography. Unlike the virtual worlds of World of Warcraft and Second Life, there is no internationally agreed-upon Grindr environment within which interactions take place; they remain rooted in the particularities of local culture, language, and interpersonal practice. The experience of using Grindr differs, at least to some extent, between New York and Tel Aviv. This approach has roots in the ethnographic work of Daniel Miller and Don Slater (2000). In their study of the internet in Trinidad, Miller and Slater argue for the maintenance of spatial and cultural specificity in internet research. They insist that accounts of networked media “cannot escape into a self-enclosed cyberian apartness” (D. Miller & Slater, 2000, p. 5) — that, in essence, we can’t meaningfully understand the internet if we imagine it to be a world entirely divorced from the cultural and material practices it is embedded within. In turn, my methodological approach treats Grindr as both a virtual field site and an anchor for a multi-sited, inescapably material examination of multiple subtly-different incarnations of the service’s basic logic.

Research ethics also play a significant part in these considerations. Helen Kennedy (2011), for example, questions whether participation by a researcher in a virtual environment can be considered ethically equivalent to sitting in a public space and taking notes. Public data posted on Twitter, she suggests, is different than data gathered in person in a physical space. In her work, she negotiates this ethical dilemma by seeking out explicit permission to use any tweets, blog posts, presentations, code, or quotes from people, even where that content was explicitly released into the public domain (for example, using something like a Creative Commons license). Other researchers, including some who have conducted participant observation on Grindr (Blackwell et al., 2014), have addressed these concerns by explicitly identifying their status as researchers in their profiles, and securing formal permission from individuals before including them in discussions of their findings.

Because of the reduced focus on individual users in this study, the risks posed by participant observation are less significant here than in other, more directly user-centric studies. Nevertheless, in my own work, I have openly identified on my profile as a researcher studying gay social media. I have not established a separate profile solely for research purposes. In this dissertation, I protect the privacy of users by excluding quotes obtained in any venue outside of a formal interview, and by omitting discussion or illustration of any personally-identifiable information contained in user profiles (including photos, even when they may be partially obscured). While the presence of other people on social networking services is undoubtedly a part of an overall user experience, the emphasis in this study is on revealing the interplay between the mechanics of software's interactive components and general patterns of user behavior,

rather than on the more sensitive actions of other individuals in their particularity.

Ethnomethodology, and data gathered through participant observation, serves to buttress and confirm analyses primarily articulated through other empirical approaches.

Textual analysis

In addition to interface analysis and on-the-ground participant engagement with these services, I conduct textual analysis of a variety of sources. For particular applications and services, these sources include policy documents (e.g. license agreements, terms of service, content guidelines), support pages, marketing materials, white papers, company blog posts, and user reviews posted on the Apple App Store and Google Play — a category of data that John Caldwell (2008) refers to as industrial “deep texts.” I examine texts published between 2009 and May 2015. I treat these industrial texts as an important source of information about the logic of the applications and services I study.

Additionally, because analysis of these documents has been an ongoing project for me since 2009, I am able to use an archive I have maintained of changes to these texts as a way to historicize the ongoing evolution of policies and practices on gay-targeted social networking services.

These textual resources offer a concrete articulation of how service providers understand their own behavior and their relationship to their users. As Kelly Gates puts it, “promotional material ... is especially important because it is here where developers negotiate the meaning of the technologies, provide a language for how to think about them, and attempt to establish their legitimacy” (Gates, 2013, p. 249). It’s critical, as Caldwell does, to treat these texts skeptically — as reflections of what industries want to

reveal about themselves, rooted in their needs as businesses pursuing customers, advertisers, and sources of funding.

I also examine trade and popular press articles about gay-targeted social networking services, as well as blog posts related to them on prominent gay-targeted blogs (such as Towleroad and Queerty). These articles offer insights into the popular reception of these services, both by gay men and by members of the public at large. Especially in the case of gay-targeted blogs, they also provide important and timely data about how individuals within gay communities respond to new features or policies. In addition to my own active, daily readership of relevant news sources and blogs, I rely on daily Google Alerts with the query “grindr” to surface texts. In practice, the number of articles that meet each of these criteria each day is small enough to allow me to examine all the texts as they are published.

Interviews

I conducted two series of interviews as part of this research. First, when appropriate and possible, I use one-on-one, open-ended interviews with designers, developers, and executives, conducted either in person, over the phone, or over e-mail, to supplement industry discourses established through textual analysis. In total, I conducted six interviews with members of the executive, development, and design staff of gay social networking services (including Grindr), of which I received permission to quote directly from one. These interviews inform my analysis throughout this work, but are especially prominent in the discussion of content management practices (chapter 3). Second, in examining user attitudes towards the data collection and targeted marketing practices of

Grindr (chapter 4), I conducted fifteen semi-structured interviews with users of Grindr who identified as having paid a monthly subscription fee for the premium version of the app on at least one occasion. While my discussion generally does not focus on the particularities of individual users' behavior, these interviews focused on how users conceptualized their choices to spend money on the Grindr service, and how that choice impacted their expectations around privacy and usability.

My use of interviews with industry figures warrants some further discussion. In practice, I treat interviews with industry figures as supplemental data sources because of the essential problems of access that emerge as a result of studying industrial contexts (Gamson, 2009; Ortner, 2010). Many businesses — and, especially, technology startups — treat interviews with members of the public (including academic researchers) as a potential source of risk for unintentionally revealing proprietary information. Few, if any, are willing to openly discuss in detail the logistical details of the processes I examine in my research (such as content management and design practices). Additionally, even when these interviews have been possible, I question the credibility of high-level industry members as informants in academic research. As John Caldwell (2008) puts it, cultural industries suffer from an “inverse credibility law” the further up an interviewer travels in an organization. My own experiences interviewing executives at startups developing gay-targeted applications support this position. Nevertheless, where possible, I sought out individuals from the companies in this study for formal and informal conversations about my findings.

Additionally, in this study I want to emphasize the importance of a data source that exists at the nexus of textual analysis and interviews: the long-form published

dialogue between the founder of a technology startup and a member of the press. These interviews, frequently published in venues like *Bloomberg BusinessWeek*, as well as on blogs like *The Verge* or the “Bits” column in *The New York Times*, offer more direct access to the thoughts and beliefs of individuals who are at the top of an organizational hierarchy within a particular company. These accounts, which present themselves as only-slightly-edited transcripts of a conversation, do not necessarily resolve the problems of source credibility highlighted by Caldwell; and, in some cases, they reflect a tendency towards hagiography in journalistic accounts of Silicon Valley culture. Despite these concerns, however, I suggest that these published interviews offer an important reflection of the values and belief structures of a software development organization.

Grounded theory

Rather than beginning with a preconceived set of characteristics I expect to observe, I rely on a grounded theory approach when examining the textual, visual, and interactive data drawn from the services and applications I examine. Following Glaser and Strauss (1967), I use grounded theory as a way to organize the diverse forms of data I gather into sets of related concepts, from which broader typologies — and, ultimately, theoretical explanations — can be derived. This approach synthesizes experiential evidence drawn from participant observation with careful textual and visual analysis in order to offer a model for why certain tendencies — in interface design, user behavior, or regulatory practice — have emerged in the real-world construction and use of gay-targeted social networking services.

A note on expertise

In October 2014, I attended the annual conference of the Association of Internet Researchers in Daegu, South Korea. On the last morning of the conference, I attended a panel with the title, “I Met My Boyfriend Online.” The panel’s participants, three gay men, each of whom had met their boyfriend (or, in one case, now-husband) online, suggested that their status as experts in the academic study of gay men’s use of networked media emerged, in part, by their own engagement and personal experiences with these services. Being a man who met his boyfriend online constituted a form of tacit knowledge (H. Collins, 1999) that could only be obtained through an especially intimate form of insider participation in gay-targeted networked media. While I’m skeptical that it is the case that only gay men can productively study gay-targeted media, the broader question of personal expertise warrants closer examination.

I bring up the personal dimensions of this panel not as a way to discredit the valuable insights the participants raised, but rather as a way into a frank discussion of my own insider status with reference to the services I study. I, too, am a man who met his boyfriend online. In fact, in the 12 years since I came out of the closet as gay, I struggle to recall more than a very small handful of dates or relationships I’ve had that *didn’t* originate online. My identity and experiences as a gay man — from coming out to a friend over instant message, to forming a relationship with someone I met on a geosocial app — are intertwined with my status as an academic researcher focusing on gay-targeted social media. I am not, to borrow from the traditional description of an ethnographer, a “professional stranger”; my insider status speaks directly to the directions this study takes.

This insider status is far from unprecedented — in studies of gay media, and in academic research more generally. Elija Cassidy’s work on gay men’s use of Gaydar and Facebook (Cassidy, 2013) begins with a highly personal narrative about the author’s own experiences using these two services to form a connection with the man who ultimately became his fiancé. His dissertation is not autoethnographic; but it — in my view productively — situates Cassidy’s interests and analysis within his everyday experiences as a gay man and a user of gay media. We can find similar reflexive positioning in several other recent studies of gay-targeted social media (Crooks, 2013; Gosine, 2007; Raj, 2011).

More generally, I draw on traditions in critical feminist theory (P. H. Collins, 2002) and feminist studies of science (Haraway, 1991; Harding, 1991) that emphasize the value of epistemologies derived from particularly situated subjectivities: of knowledge that sometimes can only be surfaced as a result of the unique experiences of individuals who are members of particular, bounded social groups. The traditional claim of this research is that a feminist standpoint has access to the “truth” about society, by virtue of its oppositional position vis-a-vis traditional masculinist viewpoints. For my part, I’m ultimately less concerned with asserting that my particular standpoint has access to the “truth” about the social situations I study, and instead would simply like to acknowledge that this study — grounded in empirical research and structured by qualitative methodologies — remains rooted in my own experiences as a gay man, with a corresponding perspective on the role of networked media in shaping the construction and expression of sexuality.

Despite this, and despite my broad recognition that a researcher is never truly outside of the texts he writes, my experiences as a user of these services is largely absent from this work. This isn't because I'm concerned that my experiences are unrepresentative (though, perhaps, they are), or that I'm somehow embarrassed about sharing my own insider status in the context of a dissertation. Rather, I want to leave myself out of this text whenever possible in service of the broader goal of leaving the particularities of individual experiences out of my analysis. My work focuses on the affordances of software, and the properties of technosocial systems that make certain outcomes more or less likely to occur, with the recognition that these systems are ultimately what users make of them. What *I*, as one user, made of Grindr isn't especially revealing of those properties; I rely, instead, on extensive textual and interface analysis to guide my interpretation of these services. My hope is that many users will find that these accounts speak to their experiences, even if their own usage has varied.

It's worth adding, as a final note, that as part of this dissertation's practical project I offer concrete design, policy, and implementation suggestions to address some of the concerns I raise about the services I study. These suggestions range from revisions to the text of an End User License Agreement to snippets of code. Where I offer these suggestions, I provide them on the basis of my experience both as an academic researcher, studying digital media design and policy, and as a practitioner, with work experience at several major technology companies in technical and policy roles. As part of my employment but separately from my academic research, I've also completed a number of courses in user experience design. Despite this experience, I'm neither a programmer nor a designer; where I provide them, my suggestions are a model of what a

solution to the problems I discuss could look like. Hopefully they will point the way toward better experiences for the millions of gay men who use the services I study.

CHAPTER 2

BIRTH

“A/S/L?”

“26/M/PA. U?”

For participants in online chat rooms in the 1990s, the three letters “A,” “S,” and “L” became a key tool for self-disclosure and identity construction. Representing “age,” “sex,” and “location” (respectively), the query, “A/S/L?” provided a consistent syntax for establishing some basic, salient personal details about the other participants in an otherwise anonymous conversation. Across a multitude of chat rooms on Internet Relay Chat (IRC), Microsoft’s MSN, Yahoo, and America Online, “A/S/L?” constituted a first step towards making the complexities of embodied identity legible in digital spaces. A/S/L empowered users to share information about themselves with the unknown others of a chat room, or to invent a new persona and en flesh her with the rudimentary outlines of an identity. The freedom to construct an identity — or multiple, fanciful identities — was only a few keystrokes away. How far — or perhaps, how not far at all — we’ve come since then.

In 2015, social networking services make considerably more complex demands for self-disclosure from their users. We’re asked to take and upload photos of ourselves, and curate ever-growing collections of tagged images submitted by family, friends, and acquaintances. We share our political affiliations and religious beliefs by typing them into form fields or selecting from an automatically-generated menu of options. We tell matchmaking websites our height, weight, hair color, and eye color using drop-down lists, and confess our turn-ons and pet peeves in terse text entry fields. Sometimes, we

even let our smartphones, tablets, and computers do the disclosing for us, sharing our current location — accurate to within tens of feet — to connect us with the people, places, and events in our vicinity. Most recently of all, the Apple Watch is able to send a real-time recording of the wearer’s heartbeat to any designated individuals (provided, of course, they have an Apple Watch, too), a feature Apple has described as a “simple and intimate” way to share an essential part of one’s self with friends and loved ones.⁷

From GPS coordinates to heartbeats, drop-down menus to selfies, there seems to be an ongoing and incessant expansion of the quantity and depth of personal information we share through digital media. Yet, faced with easy access to all this data, we need to pause and ask: What are the consequences of these new ways of sharing for how we understand each other — and, perhaps more importantly, how we understand ourselves? How do the interfaces we engage with every day structure the process of self-disclosure? What are we able to share — and what gets left out? And, how can we build and refine social software to make digitally-mediated self-expression safer and more representative of the diverse people we call “users”?

This chapter takes up these questions through an examination of the Grindr interface. Using the technique of interface analysis, I advance a framework for conceptualizing self-expression and interpersonal observation on Grindr. Herein, I focus on the interface design and interactive experience of the Grindr application, identifying how the Grindr profile creation process both puts users under external surveillance, and encourages them to carefully surveil themselves and their fellow users. This analysis builds on micro-level analysis of a visual archive of the changing design of the Grindr

⁷ <http://www.apple.com/watch/new-ways-to-connect/>

application, as well as three years of applied participant observation across a range of local contexts. I use this data to articulate a macro-level account of how Grindr's design and use create particular patterns of embodied self-expression for the service's millions of users worldwide.

As a way to systematically break down the complex interactive entity that is a mobile social application, I disaggregate the Grindr app into three discrete layers: *infrastructural*, *personal*, and *social-spatial*. A layered approach allows us to build an account of self-expression from the ground up, interrogating in turn how everything from the hardware of a mobile phone to the placement of a profile field plays a part in structuring what users are able to share about themselves. Each of these layers corresponds with particular schemas of individual, group, and institutional observation, control, and expression: of services monitoring and managing their users; of users scrutinizing themselves, and selectively sharing the results of that scrutiny; and of users observing each other and adapting their behavior on an ongoing basis.

This analysis suggests that the work of identity construction and interaction in mobile interfaces takes place both in and across these three layers. I argue that mobile social interfaces are sites of *vertically-mediated interaction*, in which multiple types and layers of data, each with different origins, affordances, and constraints, converge within a single interactive scene to create new, data-driven subjectivities. These subjectivities are the result of technological practices of observation and disclosure, and represent a new, machine-legible form of gay identity. Many of the interface conventions I discuss are unique to Grindr, or to the genre of gay-targeted geosocial networking services, and reflect the particular bodies, identities, and interactions privileged by Western gay

culture. But, more generally, this analysis speaks to a new form of networked subjectivity that we can find widely reflected across the ecosystem of social media. These networked subjectivities are at once vital, messy, and alive, even as they are deeply embedded in machinic structures of automated sorting, filtering, searching, and processing. This chapter unpacks the infrastructures, interfaces, and interactions that enable and shape their creation.

The infrastructural layer

Before we can examine the Grindr application itself, it's important to recognize the work of the devices, technologies, and relationships that enable its existence. Beneath the user interfaces of networked media are a set of infrastructural components that constitute the technical or financial preconditions for the operation of a given service or application. I identify four primary elements of the *infrastructural layer* of mobile applications: (1) Hardware and technologies; (2) Software distribution processes; and (3) Financial support and monetization. Importantly, "infrastructure" in this context does not imply neutrality (Galloway, 2004). As this section discusses, each component of the infrastructural layer reveals not only a material precondition, but also a basic set of (often tacit) normative interventions introduced by the technical or financial actors that develop or support social applications.

Hardware and technologies

The technical history of networked sociability extends far beyond the relatively recent phenomena of mobile applications like Grindr. Early sites of online sociability like dial-up bulletin board systems (Stone, 1991), The WELL (Rheingold, 2000; Turner, 2005;

2006), Internet Relay Chat (Campbell, 2004; Israeli, 1995; Wiley, 1995), and Multi-User Dungeons (Dibbell, 1991; Kendall, 1998; Turkle, 1994; 1995) offered users new opportunities to form identities, relationships, and communities in networked space. Each medium built upon the available technologies of its day to provide users with evolving capabilities to express themselves and share information. On Relay Chat (IRC), for example, text-based exchanges were the primary form of identity construction. As the IRC platform matured, users began to exchange photographs, augmenting the dynamics of information-sharing and, consequently, the networked presentation of self (Slater, 1998). The key takeaway, for the purposes of this analysis, is that certain technical properties of a medium, like bandwidth constraints or processor speed, can have important consequences for the types of interactions that medium facilitates (Stone, 1995).

Mobile social networking services like Grindr build upon two key hardware developments: (1) The popular availability of smartphones with always-on cellular internet connections; and (2) the integration of Global Positioning System (GPS) technologies into consumer devices, including, crucially, smartphones. Both developments played a crucial part in enabling the long-standing practices of networked sociability to become mobile, location-aware, and popularly accessible. Alongside this, each development also introduced new technical and commercial forms of surveillance and management, often in ways that are not clearly revealed to users.

The popular rise of smartphones — and the specific ways in which the smartphone industry developed in the mid- to late 2000s — constitutes an important foundation for how Grindr looks and works today. While devices straddling the line

between portable computers, personal digital assistants (PDAs), and mobile phones existed as early as 1996 (in the form of the Nokia 9000 Communicator), the era of mass smartphone adoption began in earnest with the release of the Apple iPhone in 2007, and the launch of the first Android smartphone (the T-Mobile G1) in 2008. These devices were the starting point for the contemporary smartphone paradigm: that is, devices that combine high-speed cellular data connectivity with large, capacitive touchscreens, coupled with easy-to-use, graphics-heavy interfaces.

This hardware and design paradigm speaks directly to the assumptions underlying Grindr. First, Grindr is built to highlight visual information: while Grindr offers users the ability to share textual and quantitative information (discussed below), it places a central emphasis on photos. The four- and five-inch displays common to many 2015-era smartphones provide a relatively limited amount of space within which apps can display information; the conventions of graphics-heavy interface design result in a tendency to highlight images, rather than text.

Second, and arguably more importantly, the fact that Grindr requires a relatively recent smartphone creates particular questions around who has access to the space of third-generation gay social networking. Grindr is available on both Apple's iOS and Google's Android platforms, a platform-agnostic approach that allows the app to reach the vast majority of smartphone users. Despite this, only relatively recent iOS and Android devices capable of running the newest versions of the iOS and Android operating systems are able to download and install the Grindr application; older and less expensive devices may be incompatible. Likewise, these services rely on always-on access to cellular data networks — capabilities that are relatively common in Western

countries but which may be prohibitively expensive or altogether unavailable elsewhere. These are broad concerns about access that are not limited to gay-targeted social networks; but they speak to what types of individuals tend to appear in the grids of profiles displayed on Grindr. The look of the Grindr Cascade is influenced by who has access to the Cascade — and that access is constrained by broader economic questions of smartphone distribution and use, both in the United States and globally.

Locative capabilities represent the second crucial foundational technology behind apps like Grindr. The ability to accurately, rapidly, and automatically determine a user's location constituted an important technological innovation with dramatic effects on the structure of mobile social networking. Even as GPS technology has now become a mainstay of devices as diverse as smartphones, tablets, digital cameras, and automotive navigation systems, its history speaks to a more complex relationship between personal data and institutional or governmental structures of surveillance. GPS was first developed in the 1970s as a military technology designed to enable precision weapons targeting (Kaplan, 2006; Parks, 2001). But by the late 1980s and early 1990s, GPS increasingly was positioned as a personal technology: a way for individuals to precisely and technically map their trajectories through geographic space. GPS, Lisa Parks writes, represents a fusion “of the intimate particularity of the personal with the broader contours of the social” (Parks, 2001, p. 219). Yet, as Caren Kaplan reminds us, the personal dimensions of GPS should not overshadow its origins as a cooperative venture between civilian, governmental, military, and commercial interests (Kaplan, 2006, p. 696) — an element of what James Der Derian (2001) calls the “military-industrial-media-entertainment network.” Geolocation, made accessible through consumer GPS

technology, has allowed a militarized system of positioning to become a seemingly intimate technology of the self.

In practice, GPS and location information function seamlessly as part of the Grindr service. Alongside data that users voluntarily share about themselves in the form of a profile, Grindr automatically collects geolocation data from users' devices any time the application is active. This geolocation data, in the form of latitudinal and longitudinal coordinates, is associated with users' profiles in order to assemble, for each individual user, a list of other nearby users — a view Grindr's developers term "the Cascade." Even as the Grindr interface does not share precise coordinates with users — location sensitivity is always presented as relative distance — the Grindr service receives and retains location information from users' devices that is as precise as the device itself is able to provide. Grindr receives and acts upon, but does not share, this highly granular level of location data.

Software distribution

The second key innovation that enabled the popularization of mobile social networking services like Grindr was the emergence of centralized mobile software distribution platforms like the Apple App Store and Google Play. These platforms give users a unified point of entry into the app economy: search, distribution, payment, and many elements of customer service are handled centrally through the app store, reducing the logistical burden on individual developers and making it easier for consumers to purchase and receive the apps they're interested in. The explosive popularity of app stores — more than 60 billion apps downloaded and \$18 billion in sales between July 2008 and October

2013, in the case of the Apple App Store (Apple, 2013b) — speaks to the appeal of this centralized business model. Via app stores, mobile software moved from the domain of the highly technically-savvy to the mainstream.

This mainstreaming tendency introduced a series of important tradeoffs around control and supervision in the app economy. Specifically, both Apple and Google enforce a specific set of guidelines regarding the types of content developers are permitted to include in their applications (Apple, 2013a; Google, n.d.). Apple goes further, individually reviewing each application submitted to the App Store before it is released to the public — a process Luis Hestres (2013) valuably recognizes as a decidedly non-neutral intervention on Apple’s part into the operation of the app economy. Both Apple’s manual review process and the developer guidelines from both Apple and Google stress that certain types of content, including types of content created by users themselves (rather than developers), are not permitted in apps distributed through their app stores. Discussions of obscene or pornographic content figure prominently in these guidelines. These restrictions on the content permitted in mainstream mobile applications has significant consequences for the types of social services users have access to. In the case of gay-targeted social networking applications, restrictions on suggestive or erotic content have had a particularly pronounced effect on the behavior of developers and service providers, often prompting them to impose strict and highly specific restrictions on user-generated content (see chapter 3).

Financial support and monetization

At the heart of these negotiations around user-generated content and personal information is a set of principally commercial concerns. Grindr proudly announces on its press information sheet that it is supported exclusively through private funding, rather than from venture capital. In practice, Grindr's revenue stream comes from in-app advertising and subscriptions to the \$11.99-per-month Grindr Xtra premium service. All the user data that Grindr has access to — from highly specific location information to demographic data gleaned from carefully managed user profiles — figures prominently into the service's financial strategies. As is the case on many networked platforms, user data becomes a commodity to advertisers, eager to leverage available information about potential consumers to more effectively target advertisements (Andrejevic, 2011; Fuchs, 2011; Langlois, 2013; Terranova, 2000; Turow, 2012; van Dijck, 2013b).

The free version of the Grindr application, available on iOS, Android, and BlackBerry smartphones, features a prominent banner advertisement across the bottom of the app's interface, visible at all times when someone is browsing the profiles of nearby users. The service's developers suggest that these in-app ads offer interested marketers a unique opportunity to reach a lucrative market of potential customers. Beneath a headline that reads, "Reach our users," an informational page of the Grindr website describes the benefits of advertising in the app:

Whether you want to target customers in your neighborhood or around the world, Grindr is the ideal way to reach highly engaged people exactly where you want to reach them. Whether you have a restaurant looking to highlight a special night or a retail business seeking new clientele nearby, Grindr is the most efficient way to talk directly to local customers. Is your business more global or internet-based? National advertising campaigns

are an efficient way to reach Grindr users across your country and around the world.⁸

This self-branding explicitly positions Grindr as distinctive in two ways: engagement and location. First, the service’s developers consistently stress how engaged Grindr users are with the app. The service’s press fact sheet suggests that the average Grindr user spends two hours per day actively using the app, on at least eight separate occasions throughout the day (Grindr, 2013b). It also indicates that Grindr users around the world exchange more than 30 million messages and 2 million photos per day. For a service with 7 million active users, this level of engagement is significant — and, from the perspective of potential marketers, is understandably appealing.

Second, Grindr emphasizes the importance of location information in enabling highly targeted marketing campaigns. This targeting takes place on every scale of data, and to increasingly varied ends. Merely by virtue of opening an application targeted at men who are interested in men, Grindr users are presorted into demographic categories that are potentially useful to advertisers. Location data offers yet more granular opportunities for targeting. This is consistent with the service’s overall emphasis on location; the service’s tagline, “zero feet away,” uses physical proximity as an index of successful interactions. But, significantly, the use of location to target advertisements suggests a new dimension to “zero feet away”: of businesses being zero feet away from their next customers. Grindr takes advantage of the highly precise information it has about its users’ positions in space — namely, their GPS coordinates, usually accurate to within several dozen feet — to offer advertisers the ability to reach engaged, active

⁸ <http://grindr.com/advertise>

customers who are nearby, in real time. This framing highlights the essential polyvalence of location data: it at once defines the localized social context of interpersonal interaction on Grindr, even as it also circulates between marketers and Grindr's ad sales representatives as a locus for effective targeted marketing.

Implicit in Grindr's suggestions to marketers are a series of demographic sorting mechanisms. Namely, by virtue of its status as a gay-targeted social networking service, Grindr is able to offer marketers access to a relatively consolidated and coherent body of users. Lisa Nakamura has described this approach, borrowing from marketers' own language, as a "tribal" advertising strategy, taking advantage of users' interface-driven self-disclosures to neatly sort individuals into easily monetizable demographic categories (Nakamura, 2002, pp. 122-123). Perhaps coincidentally, the Grindr interface also explicitly asks users to identify as members of different "Tribes." It is not difficult to imagine these demographic categories becoming yet more granular loci for marketing efforts.

Independent of other demographic characteristics, the fact that many Grindr users are gay men is an appealing trait for marketers. Katherine Sender (2004) notes that, particularly since the 1990s, gay men (and, to a lesser extent, lesbians) in the United States have been recognized as a lucrative consumer group. Motivated in part by the economic recession of the early 1990s, marketers identified three attributes of the gay market that offered the possibility of recession-proof success: education, affluence, and consolidation. What began as a halting flirtation with the gay market became, throughout the 1990s, a wholehearted embrace of the economic power of the gay consumer. Gay couples were marked by their status as "DINKS" — consumers with Double Incomes and

No Kids — and marketers seized upon the opportunity to reach a market seemingly characterized by its exceptional disposable income. The “Pink Dollar” and “Pink Pound” became a shorthand way to refer to the presumed abundance of disposable income available to gay consumers (Bengry, 2009). Edward Ingebretsen (1999) has described this emerging consciousness as the rise of “the shopping queer” — an important reframing of sexual orientation as a predictor of available disposable income and propensity to spend.⁹ Grindr takes advantage of the image of the shopping queer to position ads on the service as an efficient use of limited marketing resources.

This approach to targeted advertising works differently than many implementations of advertisements in mobile applications. In most cases, application developers exclusively contract with advertising networks like MoPub, JumpTap, Google AdSense, or Apple iAd, who in turn manage the placement of specific advertisements within the app interface. These network-based ads take advantage of information that large networks have about individual users (for instance, through the use of persistent device identifiers) to attempt to display targeted advertisements across different applications. Advertisers, in turn, buy particular ad placements through networks, based on the information those networks make available to them. The granularity of this targeting is limited by how much information — if any — is available to ad networks about the characteristics of a given app’s audience. In contrast, Grindr’s marketing offerings give advertisers the ability to take advantage of considerably more information

⁹ It’s worth remembering that this characterization of the gay market is based on a set of extremely broad generalizations: namely, that the gay consumer is (statistically speaking) well-educated and affluent. I don’t dispute that these labels are consistently borne out by demographic surveys of the American gay community; nevertheless, their pervasiveness in marketing literature has had durable effects on attempts to target gay consumers. Businesses have found themselves hard-pressed not to at least consider the gay market in their branding, behavior, and targeted advertising campaigns.

about potential ad targets than is typically available to them through mainstream online marketing channels, by virtue of the extensive customer information Grindr has about its users.

The service's developers have not hesitated to deploy this data to their financial advantage: as Grindr's CEO put it, "If an advertiser is looking to target a gay audience, there's no place better than Grindr, because we have the largest gay audience in the world" (Erlichmann, 2012). And, more recently, Grindr has begun using its access to user information to promote a political agenda, with initiatives like a get-out-the-vote campaign in 2012 and alerts about local political issues affecting LGBT individuals (Flock, 2012; Grindr, 2012). Even as they are relatively unobtrusive within the Grindr application, the presence of targeted advertisements and political messaging suggest that the technical infrastructures of social services have pronounced normative lives — albeit ones that, problematically, are only rarely made visible to users.

The personal layer

A recent and growing body of scholarship has turned its attention to the management of selfhood and identity online (Barbour & Marshall, 2012; Baym, 2010; Hongladarom, 2011; Papacharissi, 2002a; 2002b; 2009; van Dijck, 2013c; Wittkower, 2014). Using Erving Goffman's classic study of the dramaturgical presentation of self (Goffman, 1959), these works suggest that the users of networked media — and especially, of social networking services like Facebook — exercise significant individual agency managing their profiles and identities across various platforms. These studies argue that users elect to reveal or hide particular pieces of personal information in order to construct a desired

presentation of self, specific to a given social context. This performance is tantamount to a social game: a recursive process of disclosure, reception, interpretation, and reaction among actors and audiences. The appeal of Goffman's theory in networked spaces is readily apparent. But, as this section argues, dramaturgical accounts of self-presentation in networked spaces prioritize the information games of disclosure and concealment at the expense of robustly accounting for the work of the body online. On Grindr, the representation of a user's body encoded in a profile is not merely a Goffmanian performance of self; I argue that it is the result of an extensive process of self-surveillance and confession that, through software, makes physical bodies digitally legible.

Online dating has proven to be a particularly problematic case for studies of online self-presentation that rely on Goffmanian frameworks. In particular, Ellison et al (2006) emphasize that online daters find themselves caught between two seemingly opposing forces: on one hand, with the pressure to selectively highlight attributes perceived as positive or appealing in order to secure dates; and on the other, with the possibility of in-person disappointment as a foil to creative license in profile construction. The authors argue that the intimacy of the online dating's outcome differentiates it from other networked spaces: users of online dating services "desire agreement between others' online identity claims and offline identities," and are accordingly less willing to accept dissemblance, exaggeration, or selectivity (Ellison et al., 2006, p. 419). Users, they argue, view profiles as a promise, not a performance (Ellison et al., 2012).

One way that online dating services have begun to address the problem of profile dissemblance is by linking an online dating profile to another, presumably more reliable,

social network presence. The popular mobile social service Tinder, for example, exclusively draws profile information for its users from Facebook, noting that it uses data from Facebook “to make sure [users] are matched with real people.”¹⁰ The underlying logic of this strategy is authentication: As Facebook has become increasingly ubiquitous, other social services capitalize upon Facebook’s databases of personal information to connect users to their “real” identities. While this approach has significant pragmatic appeal, both for developers and users, it does not, ultimately, resolve the theoretical problems posed by dramaturgical performance; it merely defers them onto a bigger and *only presumptively* more reliable service like Facebook.

Unlike Tinder, most gay-targeted social networking services give their users the opportunity to create profiles — and therefore embodied identities — anew for themselves within the context of a particular service. This has important consequences for their users. David Gudelunas (2012b) notes that many gay men maintain, purposively, a range of different profiles — sometimes as many as seven or eight — across social networking services to allow them to pursue different social or sexual goals in different networked spaces. As a number of researchers have discussed, this careful management of identity across different online platforms gives users the opportunity to selectively reveal information in accordance with their individual perceptions of the purpose or norms of a given service (Papacharissi, 2002a; 2009; van Dijck, 2013c) or expectations of their audience (Marwick & boyd, 2011). Embedded in this process of selective revelation is more than just Goffmanian performance; as I discuss, users surveil both

¹⁰ <http://www.gotinder.com/faq/>

themselves and the other users of a given platform in order to negotiate these boundaries of self-disclosure in order to create a proper platform-specific subjectivity.

Upon launching the Grindr application for the first time, users are presented with a view of the Cascade showing other profiles in their vicinity, along with a blank, gray square in the upper-left corner, representing the user's own profile, yet to be created. Even as users are not required to disclose any information about themselves in order to browse the Cascade or the profiles of other users, Grindr does not permit users with entirely blank profiles to communicate directly with other users. The visible profiles of other users serve as a source of encouragement for new users to create a profile and, in turn, become visible themselves. Voyeurism, while permissible, does not grant access to the interactive core of the Grindr service. Creating a profile, therefore, is a central part of how users engage with Grindr.

When creating a profile, users disclose three types of information: (1) photos; (2) free-form text; and (3) quantifiable or categorical data about themselves or their interests. I argue that this process of constructing a profile from an agglomeration of data types constitutes an important moment of self-surveillance. Users are encouraged by the Grindr interface to consider their bodies and identities when deciding how to represent themselves on the service. Each type of data solicited by the Grindr profile creation process brings with it a corresponding set of surveillant and classificatory practices which warrant more specific theoretical positioning.

Paulo Vaz and Fernanda Bruno (2003) provide a constructive point of departure for this positioning in their discussion of self-surveillance. Even as academic studies of surveillance have fixated on Michel Foucault's (1978a) discussion of the panopticon, Vaz

and Bruno suggest that the normalizing discourses of care, risk, and abnormality offer an important locus for understanding the internal operation of panoptic power. Panoptic surveillance, they argue, relies on both the potential supervision of the panopticon (the traditional interpretation of Foucault's argument in *Discipline and Punish*), as well as on an internalized conception of normal subjectivity that prompts individuals to monitor themselves in ways that correspond with social expectations. Self-surveillance, Vaz and Bruno note,

is usually understood as the attention one pays to one's behavior when facing the actuality or virtuality of an immediate or mediated observation by others whose opinion he or she deems as relevant — usually, observers of the same or superior social position. [W]e propose to open the concept to include individuals' attention to their actions and thoughts when constituting themselves as subjects of their conduct. (Vaz & Bruno, 2003, p. 273)

Self-surveillance, in this definition, becomes critically associated with the care of the self (Foucault, 1986), in which care “assumes the form of an effort to constitute oneself as a normal citizen” (Vaz & Bruno, 2003, p. 279). I argue that this negotiation of normality in the presentation of self is deeply embedded in the structure of social networking services.

On Grindr, self-surveillance first takes place around the production and selection of the information displayed on a user's profile. The profile photo, subject to Grindr's content management policies (see chapter 3), is the focal point of a user's profile. Occupying the entire interface when viewing a profile, the single profile photo constitutes a critical point of identity construction on Grindr. Beyond a photo, users can add a range of textual, quantitative, and categorical information to their profile. The space for text on the profile is fairly limited. Users can share a display name, a headline, and a 255-character free-form description labeled

“About Me” — considerably less textual information than is displayed on profiles of other social services, including dating services like OkCupid or Match that also have mobile applications. Additionally, users can choose to share their:

- Age
- Height




Fig. 1: Grindr profile creation screen.

- Weight¹¹
- Ethnicity
- Relationship status
- Body type, such as “Toned,” “Average,” “Large,” “Muscular,” “Slim,” and “Stocky”
- “Grindr Tribe” identification, which offers twelve groups that users can choose to list on their profile, including “Clean-Cut,” “Daddy,” “Geek,” “Trans,” and “Poz” [HIV-positive]
- and the type of interactions they’re “looking for,” such as “Chat,” “Friends,” or “Networking.”

Grindr offers users a range of options for self-disclosure, ranging from the explicit and quantitative to categories whose precise meaning is left open to individual interpretation. Each data point, when combined with Grindr’s extensive filtering tools (discussed in “The social-spatial layer,” below), gives users a way to both represent themselves and seek out partners who fit their requirements.

Perhaps the most prominent and widely discussed component of a Grindr profile is the photo. Many popular accounts of the service emphasize the politics of the profile photo, epitomized in a tension between the so-called “faces” and “torsos” (Habib, 2013;

¹¹ In March 2015, Grindr released an update to their iOS client application introducing several new features and a number of miscellaneous bug fixes. Shortly after the update’s release, a number of users posted negative reviews of the application on the Apple App Store, noting that it had introduced several frustrating bugs, including an apparent inability to list one’s weight as 145lbs. (Apparently, the app would instead list their weight as 144lbs.) This admittedly minor glitch in the Grindr app points to two compelling insights: First, we should remember that even quantitative or automatically-collected information is subject to error. Even as the structure of the Grindr application puts a great deal of emphasis on “objective” measures of bodies and identities as ways to organize users, there’s no guarantee that these measures are necessarily less subject to bias than any others. Second, bugs like this one allude to a branch of research that I leave largely unexplored in this dissertation but which is certainly worthy of further examination: namely, the role of bugs, glitches, and errors in structuring users’ experiences of networked media.

Kapp, 2011) — that is, between users whose profile photos show their face and those that do not (often, instead, displaying a shirtless torso). These conventionalized portrayals of self have become a type of visual shorthand within the Grindr service, pointing indexically towards the type of interaction a user is looking for. “Torsos” are presumed to exclusively seek casual sexual encounters, while “faces” leave open the possibility of other types of social interactions. Nothing in the process of profile creation makes this tacit language of the profile photo evident to users; instead, this knowledge emerges organically through the use of the Grindr service.

This approach to reading photos for information about the individuals they portray is not unique to social networking services. Crucially, the study of visual criminology, and closely related inquiries into historical and contemporary medical imaging, provide a robust theoretical framework for understanding the power dynamics embedded within the production, reception, and interpretation of photos in the context of embodiment and identity construction. Allan Sekula (1986), charting the criminological techniques of Alphonse Bertillon and Francis Galton in 1880s France, identifies photographic techniques as being at the heart of “positivist attempts to define and regulate social difference” (Sekula, 1986, p. 19). Of particular interest here is Galton’s approach, which used composite photography to attempt to discern the general visual characteristics of criminal types. These images, and other forms of visual evidence, became central techniques of institutional surveillance when they emerged in the second half of the nineteenth century (Tagg, 1988). And, importantly, the criminological impulse to view images as an index of real tendencies or events has not subsided; it has, instead,

taken on the increasingly technologically sophisticated forms of video, facial recognition, and biometrics (Gates, 2011; 2013).

Importantly, a similar analytic regime emerged in medicine in parallel with criminology. The use of the photographic image in medicine, notes Sander Gilman (1982; 1989; 1995), became a way to encapsulate the realities of a disease or patient condition in a manner perceived by practitioners as relatively unproblematic. And, as Lisa Cartwright (1995) recognizes, the emergence of new imaging technologies and techniques (like the microscope or X-ray) enabled the construction of new regimes of medical knowledge based on access to the body granted by these images. Bodily imaging techniques, Cartwright argues, became crucial instruments

in the emergence of a distinctly modernist mode of representation in Western scientific and public culture — a mode geared to the temporal and spatial decomposition and recomposition of bodies as dynamic fields of action in need of regulation and control. (Cartwright, 1995 p.xi)

Medical practices of surveillance and analysis rely on images (photographic, radiographic, and cinematographic) of the body to identify points of difference — and therefore, of pathology. In each case, and particularly in the case of radiographic imagery, bodies are subject to the suspicious gaze of medical practitioners, who use images to discern and subsequently authenticate diagnoses of illness or dysfunction.

On social networking services, the suspicious gaze is both internal and external. Users surveil each other, appraising photos and, in the event of a face-to-face meeting, assessing the correspondence between a photo and an en fleshed visage. But, critically, users are encouraged to turn a suspicious eye to their own images, selecting profile photos that not only correspond with Grindr's imposed content restrictions but which *also*

meet the dual and occasionally opposing requirements of seductive appeal and “real life” authenticity. By allowing only a single photo on profiles, Grindr prompts its users to ask themselves: Is *this* the photo that captures what I want to express about myself and my body? Does it meet the requirements of this space? Is it representative of *me*? Is it the ‘me’ I want to present *here and now*? Does it serve my objectives in this space? Answering these questions demands ongoing, adaptive, and context-specific self-surveillance. The profile photo represents one outcome of that process.

Categorical data types, such as Grindr Tribe and body type, require yet more acute form of self-surveillance. In order to select an appropriate Tribe, for example, users need to reflect upon their bodies and identities, and subsequently classify them in terms set forth by an external entity (Grindr’s developers). In contrast to the relative freedom of text entry fields or profile photo selection, categorical classification imposes strict constraints on how users can share information about their bodies, identities, and desires. This process of classification is a key part of the construction of the Grindr profile.

The act of establishing a set of twelve categories constitutes an important intervention into the expression of bodies on Grindr. As Pierre Bourdieu notes, “The fate of groups is bound up with the words that designate them” (Bourdieu, 1984, pp. 480-481). By including certain body types or Tribes as options, Grindr’s developers reify them as units of self-expression and comparison. These terms are artificial, but through a recursive process of definition and reinforcement, they acquire operational meaning for users. They become, in Ian Hacking’s terms, ways of “making up people” (Hacking, 2006).

The design of the Grindr interface is itself significant in this process. In place of the free-form text entry fields deployed elsewhere in the profile creation process, categorical data is inputted by users through structured lists and check boxes. A user is presented with a finite list of options, of which he is then given the ability to select one (or, in the case of Tribes, several) to represent him. These menu-driven identities, as Lisa Nakamura describes them, position users

within the paradigm of the “clickable box” — one box among many on the menu of identity choices. When users are given no choice other than to select the “race” or “ethnicity” to which they belong, and are given no means to define or modify the terms or categories available to them, then identities that do not appear on the menu are essentially foreclosed on and erased. (Nakamura, 2002, pp. 101-102)

Nakamura recognizes, importantly, that these constrained choices were the product of interface design processes designed to make the complexities of networked spaces more accessible to users. The anxiety and indeterminacy of the free-form text field is ameliorated by creating a finite set of possible inputs. And, particularly in the case of mobile applications, the relatively small displays of portable devices like smartphones put a premium on concise design practices that (claim to) express as much information in as little space as possible. Nevertheless, the abridged self-expression of categorical data entry reduces the ability of users to present themselves on social networking services in their own terms — whatever those terms might be. Expressing bodies and identities through the drop-down menus, check boxes, and numerical entry fields of the Grindr interface requires an act of translation: between the phenomenological complexities of embodiment and the digitally-legible expression of bodies in databases and mobile interfaces.

I want to conceptualize this practice of self-disclosure as a form of confession: of individuals ritualistically revealing details of their bodies and activities in a prescribed format. Foucault discusses precisely this type of revelation in *The History of Sexuality* (Foucault, 1978b). Ritual confession — what Foucault calls the “confession of the flesh” — became a medium concerned with the minutiae of carnal conduct. Sex, Foucault writes, “must not be named imprudently, but its aspects, its correlations, and its effects must be pursued down to their slenderest ramifications ... everything had to be told” (Foucault, 1978b, p. 19). Confession does not merely describe the recitation of the contours of an individual’s body or activities; it rather entails a careful scrutiny of their particularities, and an exacting rendering into discourse of those details. The confession — that is, an individual’s recitation of exacting details about himself — has become, in Foucault’s estimation, “one of the West’s most highly valued techniques for producing truth” (Foucault, 1978b, p. 59)

I see the essential elements of this ritualistic confession in the structure of the Grindr profile creation process. Software interfaces serve the function of providing an orderly outlet within which this confession of self can take place; they provide the framework within which individuals can make sense of and recount the particularities of their identities and desires. The task of the individual is to provide the raw material — the data — of himself; software is responsible for making this information useful: for revealing the truth of the self in the Grindr cascade.

The stakes for this process of translation through confession are considerable. Gilles Deleuze (1992) famously coined the term “dividual” to describe the endlessly subdividable subjectivities created through systems of control. The profile creation process is

one such system of control. In the context of digital technologies, writes John Cheney-Lippold,

dividuals can be seen as those data that are aggregated to form unified subjects, of connecting dividual parts through arbitrary closures at the moment of compilation of a computer program or at the result of a database query. (Cheney-Lippold, 2011, p. 169)

This new, data-driven subjectivity stands in stark contrast to traditional liberal conceptions of individuals as rational, autonomous subjects. Subjects are, as David Poster puts it, interpellated *by* databases (Poster, 1995). Through information that they either voluntarily disclose or which is passively, automatically collected about them, an aggregation of data points stands in for a fully autonomous act of declaration of self.

Identity becomes something that happens to the individual, rather than an expression of subjectivity that individuals themselves agentially author.

Returning to Sekula's discussion of nineteenth century criminology, we can see echoes of Bertillon's taxonomic approach to identifying deviance in attempts to render subjectivities legible through statistics, databases, and electronic interfaces. Bertillon,




Fig. 2: A completed Grindr profile. Grindr stock screenshot, available at <http://grindr.com/press>.

Sekula notes, “sought to break the professional criminal's mastery of disguises, false identities, multiple biographies, and alibis” (Sekula, 1986, p. 27). Cataloging and statistical analysis were key parts of this process. For Bertillon,

the mastery of the criminal body necessitated a massive campaign of inscription, a transformation of the body's signs into a text, a text that pared verbal description down to a denotative shorthand, which was then linked to a numerical series. Thus Bertillon arrested the criminal body, determined its identity as a body that had already been defined as criminal, by means that subordinated the image — which remained necessary but insufficient — to verbal text and numerical series. (Sekula, 1986, p. 33)

We can thus position the profile photo against the quantitative and categorical data entered alongside it during the profile creation process. While central, the profile image is both necessary and insufficient in the constitution of subjects on Grindr. The profile photo is anchored by the other types of data disclosed in a user’s profile: height, weight, or categorical identifications authenticate, qualify, and elaborate upon the information revealed in the photo. The Grindr-using subject can only be created through the integration of these disparate datum.

Ultimately, the process of profile creation constitutes a form of digital epidermalization — a practice of subjective construction that prioritizes superficially accessible, classifiable, and generalizable details. Simone Browne (2010), writing about the use of biometric data in legal processes like border control, notes that presumptively neutral technological practices like fingerprinting have pronounced and non-neutral consequences for particular racial or ethnic groups. Building on the work of Frantz Fanon (1967), Browne argues that digital epidermalization constitutes a

...moment of fracture of the body from its humanness, refracted into a new subject position (‘Look, a Negro!’, or an ‘illegal alien’, or some other negatively racialized subject position). It is the interpellating gaze of the

moment of contact that produces these moments of fracture for the racial Other, indeed making and marking one as racial Other, experiencing its 'being through others.' (Browne, 2010, p. 134)

The consequences for this epidermal thinking need not be limited to race. Instead, Browne highlights the ways in which digital technologies of quantification, classification, and self-expression can systematically alienate individuals from their bodies and identities. Something of the self is represented in fingerprints — or in a Grindr Tribe identification or profile photo. But something important is also lost: the intricacies and specificities of bodies and identities that cannot easily be expressed through digital interfaces. Through regimes of self-surveillance, users themselves are enrolled in this process of generalization; building a profile requires users to interpret themselves in the terms of a social networking service.

The normative consequences for this reconceptualization of self remain open to contestation. For my part, I want to emphasize the constructive potential of these new ways of understanding identities, bodies, and communities. As I discuss below, interpreting the self in the language of social networking services offers individuals unprecedented possibilities for making sense of complex social landscapes, and for finding and pursuing relationships with other people. Undoubtedly, as Browne, Fanon, Deleuze, Poster, and others have suggested, significant detail is lost in this process; but we should stop short of dismissing these practices of technosocial self-expression as inevitably resulting in a data-driven impoverishment of the self.

The social-spatial layer

Having established the infrastructure for a social service and examined the processes by which individuals can express their bodies and identities through digital interfaces, we come at last to the social. What of the other people in the Cascade? What are the practices of social observation and surveillance engendered by the Grindr interface? Borrowing from Jonathan Crary, I suggest that Grindr users are positioned at once in three observational modes: “a spectator, a subject of empirical research and observation, and an element of machine production” (Crary, 1988, p. 20). These three modes are configured vertically within the Grindr interface, consolidating a range of different spectatorial modes and techniques within a single application. As we’ve seen, the technical structures of social networking services both enable top-down surveillance and prompt introspective practices of self-surveillance. By expressing bodies-in-physical-space in the grid format of the Cascade, I argue that the Grindr interface enables a third form of surveillance: a vertically-mediated, data-driven regime of users surveilling each other.

One branch of surveillance scholarship has labeled these behaviors *lateral* (Andrejevic, 2005), *social* (Joinson, 2008; Tokunaga, 2011), or *participatory* (Albrechtslund, 2008) forms of surveillance. Mark Andrejevic, for example, characterizes lateral surveillance as techniques of peer-to-peer monitoring ranging from “casually Googling a new acquaintance to purchasing keystroke monitoring software, surveillance cameras, or even portable lie detectors” (Andrejevic, 2005, pp. 488-498). We have become so skeptical of each other, Andrejevic writes, that we turn to the technologies of surveillance to “appeal to the evidence of one’s eyes rather than the words of others” (Andrejevic, 2005, p. 482) — in short, mitigating the risk of increasingly spatially and

temporally distanced forms of communication (of which online dating is an important example) by increasing the amount of information at our disposal about the people we're communicating with.

Alice Marwick (2012) makes similar claims about practices of social surveillance that emerge on networked platforms like Facebook. We turn to behaviors like “Facebook stalking” — scrutinizing the digital output of our acquaintances, friends, or lovers — to obtain otherwise unavailable forms of social knowledge. Valuably, Marwick differentiates social surveillance from Andrejevic's discussion of lateral surveillance by highlighting that, on social platforms, *being seen* is as analytically and behaviorally significant as *seeing others*. Social media, Marwick writes, “has a dual nature in which information is both consumed and produced, which creates a symmetrical model of surveillance in which watchers expect, and desire, to be watched” (2012, p. 380). Unlike many accounts of networked surveillance, which stress top-down, institutional practices of observation and data-collection without recognizing the motivations of users in sharing their personal information in the first place, Marwick acknowledges that, in many instances, users actively elect to both share and consume personal information. A framework based on social surveillance recognizes that users themselves must buy in to social systems in order for surveillant practices to emerge — and that, consequently, users are intimately involved in various practices of surveillance.

It's important to recognize, as Anders Albrechtslund (2008) does, that these practices of networked surveillance can also serve an empowering or subjectivity-building role. Networked surveillance, Albrechtslund writes, need not reduce a person under surveillance to a passive, powerless subject; rather, being looked at “can be part of

the *building* of subjectivity and of making sense in the lifeworld” (Albrechtslund, 2008). Drawing on the example of webcams and personal broadcasting, Albrechtslund suggests that visibility can become “a tool of power that can be used to rebel against the shame associated with not being private about certain things” — making one’s self visible within a networked surveillance apparatus is, at least in part, an assertion of individual agency.

I want to emphasize an analytic framework that draws from all three approaches. Grindr makes available to its users a toolkit that enables and encourages techniques of lateral and participatory surveillance, allowing users to rapidly gather, sort, and act upon data about each other. On one level, this serves a credentialing function, enabling users to quickly synthesize visual, textual, and quantitative data contained in a profile to establish a sense of that profile’s plausibility. But, equally significantly, the Grindr service is built around the construction of networked subjectivities for the purpose of facilitating desirable social interactions, both online and offline. Users share information about themselves, and consume the shared information of others, in order to pursue any of a wide range of goals — from meeting sexual partners to exchanging restaurant recommendations in unfamiliar locales. The visibility engendered by the Grindr service is normatively complex, inflected with questions about how embodied difference is manifested through constrained digital interfaces. Nevertheless, I want to stress that the surveillant practices enabled by the Grindr service should be conceptualized with these ambivalences at heart: at the union of skeptical monitoring and subjectivity-building.

Location is a key component of both of these processes. Surveillance on Grindr takes place at the nexus of digital representation and physical proximity. Unlike John

Urry’s (2002) account virtual “co-presence” or Christian Licoppe’s (2004) discussion of telephonic “connected presence,” interactions on Grindr are anchored by the centrality of location data in the service’s operation. A profile should be authentic, because false information can easily be revealed by a passing glance from another user, potentially only feet away. And ultimately, intentional physical co-presence is a chief outcome of interactions on Grindr; as Grindr’s developers put it in marketing copy on their website, “0 feet away: Our mission for you.” Enabling the “zero feet away” interaction — that is to




Fig. 3: The Grindr Cascade.

say, physical co-presence — is at the core of how Grindr’s interface sorts and presents data about users.

Upon launching the Grindr application, the service assembles a view of the Cascade comprised of a list of nearby profiles in ascending order by distance. This “Grindrscope,” as Roderick Crooks (2013) has labeled it, is an “astonishingly local” and “ad-hoc” social space — that is to say, the Grindr Cascade dynamically reassembles itself based on a user’s present location, displaying, reordering, and hiding users based on passively-and omnipresently-gathered GPS location data. This dynamically constructed

grid interface, when it was introduced by Grindr in 2009, was a unique and innovative way to represent users in space. Unlike many social networks that use location data, such as foursquare, Facebook, and Instagram, which allow users to augment digital representations of physical space through their networked behaviors (de Souza e Silva, 2006; de Souza e Silva & Frith, 2010b; 2012; Frith, 2013; Lee Humphreys, 2007; R. Schwartz & Halegoua, 2014; Sutko & de Souza e Silva, 2011), Grindr deploys location data as a way to construct a social, rather than spatial, map of its users. Location data is used to determine relative geographic proximity; landmarks, buildings, businesses, or even neighborhood or city distinctions are discarded in the Grindr interface.

The use of location data to construct the Grindr Cascade enables a wide array of user behaviors based on both access to and the sharing of location information. For example, a user visiting a new locale may solicit recommendations for gay-friendly bars from other users in his vicinity, transforming the Grindr Cascade into an easily accessible source of local knowledge and expertise. In tandem, users also broadcast their current locations, reciprocally positioning them in the Cascade of others as a potential target for interaction.

But the constitution of the Grindr Cascade is not as neutrally or consistently location-based as Crooks suggests. The first profiles displayed in a user’s Cascade are profiles that have already been marked as favorites and profiles from which a user has a new message — stressing ongoing social ties over the assemblage of spatially proximate profiles that makes up the rest of the Cascade. And, importantly, the Grindr service provides a variety of techniques for excluding unwanted profiles from view on the Cascade. For example, users may individually block other users, prohibiting all interaction with them and removing their profiles from the Cascade regardless of proximity.

More sweepingly, users can deploy data gathered from profiles to construct custom Cascades that automatically eliminate whole categories of profiles from view. Users can elect, for example, to view a Cascade comprised entirely of profiles identified with a particular Grindr Tribe or ethnicity, or which fall within a specified range of ages. This automated filtering dramatically increases the material force of the designations disclosed during the profile creation process. Semantically ambiguous classifications like

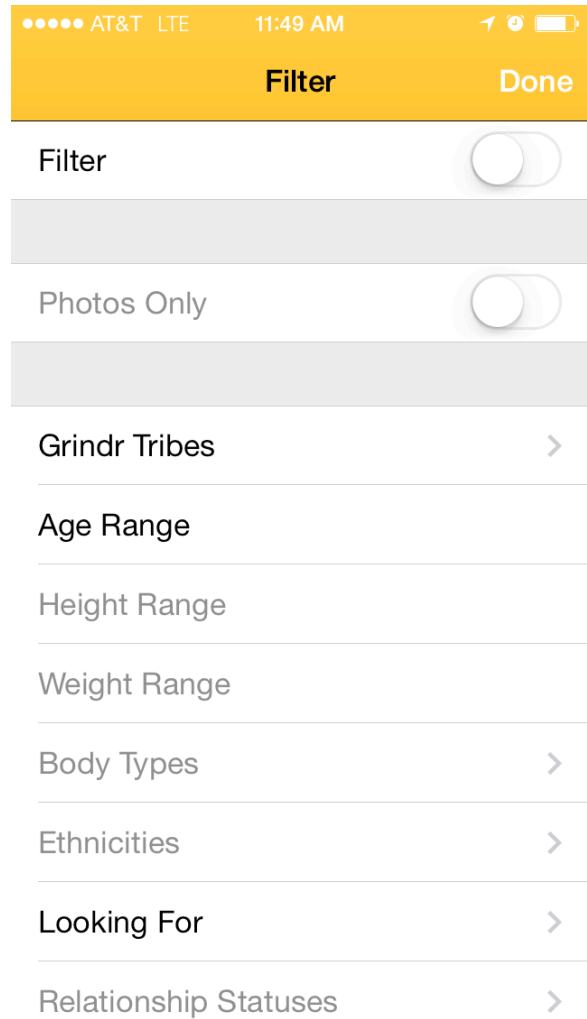


Fig. 4: Filters available for the Grindr Cascade. Filters in gray are only available for paying Grindr Xtra subscribers.

body type and Tribe become loci for social surveillance, and, by extension, action on the part of users. Further, electing not to include a Tribe, ethnicity, or body type identification on a profile means exclusion from any searches or filters that select on the basis of those classifications — in other words, a diminished social presence that threatens the Grindr service’s basic interactive objective of reciprocal visibility (Bucher, 2012). This is not, as Crooks argues, another instantiation of long-standing offline gay practices of coded disclosure of desire, such as the “hanky code.” Rather, the ability to rapidly and automatically sort and act upon the data encoded on user profiles represents a radically new and algorithmically-driven model of networked sociability, reliant upon self-surveillance and the reciprocal disclosure of machine-readable data about users’ bodies and identities. Grindr is not simply — or even principally — a digital version of analog practices of cruising or sociability; at its core is a user-friendly implementation of what Oscar Gandy (1993) has called the “panoptic sort,” introduced into the intimate sphere of personal and romantic relationships.

The user-friendliness of this implementation is at the heart of the popular success of social networking services like Grindr. The Grindr interface gathers a vast amount of data from its users and translates them into the easily-accessible forms of the Cascades and the individual user profile. I argue that this act of translation — undertaken in a wide variety of ways by virtually all social networking services — represents the vertically-mediated integration of user data into an application built around practices of participatory surveillance.

In this discussion of verticality, I draw on two distinct, but related, meanings of the word: verticality as vertical distance; and verticality as multidimensional integration. Each warrants more specific examination.

First, verticality alludes to looking down onto something — what Donna Haraway has called the “god trick” of a projected, omniscient gaze from high above (Haraway, 1991, p. 189). In this sense, Grindr gives its users an overhead view from which to watch each other, harkening to video feeds from closed-circuit television cameras, screened from afar (Koskela, 2002; Walby, 2005), or to the view of New York City from atop the World Trade Center famously described by Michel de Certeau (1988). The richness and messiness of identities and bodies are flattened and rationalized by vertical media; they are made quantifiable, sortable, and machine-readable. This flattening creates what David Lyon has termed an “electronic superpanopticon” —a digitally-enhanced all-seeing eye (Lyon, 2001, pp. 108-109). And, through the Grindr interface, this superpanoptic perspective is made readily available to users. Each photo thumbnail in the Cascade can be tapped and expanded into the full profile of a user, then returned effortlessly to its original position or banished with a tap of the “Block” button. Profiles can be sorted, filtered, hidden, and evaluated in order to reveal only the users who are most desirable. Users look down from above onto each other, taking advantage of the wealth of data available to them through a system that could not exist without their willing disclosure of personal information.

But verticality need not be equated with reductiveness. Instead, I suggest a second reading: of vertical mediation as the integration of a wide range of data types and points into a single interactive scene. This interpretation of verticality has, traditionally, been

associated with geopolitics (Graham, 2004; Scott, 1999). Eyal Weizman's influential work on the vertical politics of disputed territory in the West Bank, for example, stresses that land can no longer reasonably be understood as a two-dimensional surface; rather, it should be conceptualized as a "large three-dimensional volume," layered with cultural, political, and strategic data (Weizman, 2002). Peter Adey has made similar claims about what he calls "aerial subjects" — those subjectivities constituted by "regulations, technologies, practices, forces, and affects" that bear down upon them from above (Adey, 2010, p. 206). Crucially, however, the aerial subject is not merely a superficial projection; instead, through interpellation from above, he "has been given depth, with a surface and an interior; intentions and desires may be read or incited by an address of superficial or far more vital capacities" (Adey, 2010, pp. 206-207). Aerial subjects react to being surveilled. Their behaviors and self-presentations change based on the systems that observe and interpellate them. Aerial subjects are multidimensional in tandem with vertical systems of observation and control, not in spite of them.

The subject on Grindr, too, is multidimensional, even as the chief mode of access to him is a vertical interface that flattens a wide array of data into a grid of photos. On Grindr, vertical mediation collapses the boundaries between different aspects of the user — information about their body, personality, preferences, and location — and positions this reintegrated individual in a grid of other integrated user presences. This reintegrated subjectivity is not merely a Deleuzian individual; even as the profiles of Grindr users are constrained by technical systems, they are not wholly constituted by them. The subjectivities of the Cascade have depth, emerging from processes of both self-surveillance and reciprocal surveillance of other Grindr users. The vertical interface of

the Cascade provides users with the tools to actively negotiate the rich array of data about other users available to them within the Grindr service.

These flattened (but not altogether flat) subjectivities are part of a broader way of looking at other people — what I've termed the social-spatial layer of the Grindr app. Grindr positions individual profiles, enfolded in varying degrees of quantitative, qualitative, and visual detail, within a grid of other profiles. This results in a dramatically different interpersonal experience than, for example, browsing other users on an app like Tinder. Like Grindr, Tinder aggregates information about users from a variety of different sources, including geolocation and data gathered automatically from Facebook. But, unlike Grindr, Tinder presents profiles linearly, as a succession of individual presences. Profiles on Tinder can only ever be viewed one by one; they cannot be compared, reorganized, or casually browsed. Where Tinder puts the deepest, most detailed version of an individual's profile at the forefront of the app's interface, Grindr presents its users with a wide range of thumbnails, each of which can be (but isn't necessarily) expanded into a richer view of another person. Tinder's central interface convention — swiping left or right to indicate whether or not one is interested in the user displayed on the screen — emphasizes split-second, intuitive judgements as a way to establish connections between users; decisions cannot be undone, and profiles cannot be browsed any way other than in a single, successive stream. Verticality on Grindr is a fundamentally dissimilar way of conceptualizing the experience of looking and being looked at in a social networking service. It allows users to move effortlessly between different scales of data: from the zoomed-out Cascade to the detailed view of a particular profile. And, by giving users the power to filter, sort, and search the Cascade, individuals

are able to make this zoomed-out view of their fellow users better suited to their interests and desires. Subjectivities on Grindr are flattened; but users are fundamentally empowered by precisely this flatness.

Conclusions: Asymptotically approaching embodiment

This chapter offered a new, post-Goffmanian account of networked sociability and identity construction through a discussion of what I term vertical interfaces. I want to stress that the account of verticality I have outlined here need not be equated with the reduction of individuals to squares in a closed-circuit television system. Vertical mediation can be robust, enfolded, and multidimensional, even as it is constrained and carefully managed. The stakes for this verticality are significant, both theoretically and in practice. Users do not perform their identities through these vertical interfaces; or at least, they do not perform them in the unified manner of Goffman's front-stage acting. Instead, people reveal, selectively, under constraint, and through a process of self-surveillance, discrete pieces of information about themselves: a photo, a height, a body type, a Tribe. Technical infrastructures, both seen and unseen, operate on our information, making people legible to software — and, in the process, to advertisers. Ultimately, individuals are reintegrated and represented to ourselves and to others as a profile in a cascade of other profiles. Users scrutinize each other, and, at the point where these layers of data are integrated with each other, interaction becomes possible.

But these questions are more than just theoretical elaborations on how we understand identity work online; they can and should inform the design and implementation of social identity creation processes across a wide range of networked

platforms. The recent and widely-discussed case of Facebook’s revisions to the gender field on users’ profiles provides a clear illustration of the stakes for these discussions. Since the site’s launch in 2004, users had the option of identifying as either male or female, or choosing not to display a gender on their profile at all. In 2014, Facebook collaborated with a group of advocacy organizations to roll out a change to the gender field (ReadWrite Editors, 2014). The updated gender field offered users 56 options, including “intersex,” “gender variant,” “transmasculine,” and “two-spirit,” among many others. A year later, Facebook announced a further revision to the feature, allowing users to enter text freely into the gender field to create up to ten custom “gender terms” to describe their identification (in addition to any combination of the preexisting 56 terms). The Facebook Diversity page hailed the move as “[giving] people the ability to express themselves in an authentic way.”¹² The underlying design revision was fairly simple: changing the older drop-down menu of binary gender choices to a free-form text field with a selection of pre-populated but not mandatory options. The conceptual shift, however, is significant: Facebook introduced some additional complexity to the gender profile field in order to allow users greater freedom to construct an identity that is better representative of — or, perhaps, a more authentic version of — their lived experience.

Other services, including the romantically-focused website OkCupid, have undertaken similar projects to elaborate upon both the specific interface options available to users, and the conceptual frameworks underlying the tools sites give to users to express themselves. In the case of OkCupid, for example, the service expanded its gender and sexuality options in tandem, to better account for the ways in which sexual practice and

¹² <https://www.facebook.com/facebookdiversity/posts/774221582674346>

gender identification can be co-constitutive (North, 2014). These ideas, of course, have long had resonance in queer theory (Halberstam, 1998); but OkCupid has engaged with them at the level of design and technical practice. For some users, this resolved a “broken window” on OkCupid around the expression of transgender bodies and identities that required tactical workarounds. Historically, trans users of the platform noted that they were forced to include lines like, “Hi, I’m a trans guy” in the body of their profiles to alert visitors to their identities; in many cases, users reported that those lines went unseen or ignored, resulting in potentially dangerous, threatening, or uncomfortable matches with users uninterested in a connection with a transperson. By incorporating trans identification explicitly into the implementation of user profiles, trans users are able to make their bodies legible online, using their own terms, in ways that correspond with the expected presentation of personal information on the OkCupid platform. The result, in users’ experience, has been a service that’s better able to represent non-cis, non-hetero/homosexual identities within the established conventions of a social networking profile.

These features exist at the nexus of three sometimes conflicting tendencies: (1) ease of use; (2) freedom of expression; and (3) commercial legibility. A binary gender field maximizes ease of use: users have two simple options to choose from. It also generates straightforward information about the gender identification of users that Facebook is able to employ for targeted marketing; “male” and “female” are gender categories with clear meaning to marketers. On the other hand, a binary gender field constrains the possibilities available to individuals for expressing their identities outside of a clickable box: gender identities which don’t fit the male/female binary can only be

represented as an abstention to identify. By contrast, the free-form gender field offers users nearly limitless possibilities for sharing information about their gender identities; not only can they type in custom terms, but they can use a combination of up to ten terms to make their identity both legible and detailed. Undoubtedly, these changes make it more possible to share non-normative gender identities on Facebook; but they also increase the complexity of the gender profile field, requiring users to type their gender identity into being, rather than selecting from among a small set of prepopulated options. The changes also make it more challenging to neatly parse users into gendered categories for marketing purposes; when users can enter text freely, categorical sorting becomes less and less possible. Facebook's revisions to the gender feature prioritized freedom of expression over ease of use and commercial legibility.

These negotiations look different on Grindr. The relatively small screen sizes of mobile devices create a significant disincentive to increasing the complexity of interface features; simple, tappable boxes make the profile creation process quick and straightforward. Replacing a "body type" field with fewer than ten choices with a free-form text entry field would significantly increase the effort required to share information about one's body. It would also limit the usefulness of the filtering and search features that enable users to quickly navigate the Grindr Cascade in the pursuit of those profiles that are most likely to be a match.

Despite these concerns, I recommend an approach to profile design that minimizes the amount of pre-imposed structure in the profile creation process. In place of a clickable box or drop-down menu, one might imagine an open text entry field, wherein users can express, in their own terms, the characteristics of their bodies and identities that

are most salient to them. Absent the structure of the current profile creation process, the act of confession at the heart of this process would be one defined by an individual's own sense of what's most important to share in the limited space of a profile. We can, reasonably, expect that these confessions will contain largely the same pieces of information as the current profile: height, weight, age, and body type remain salient characteristics in the context of an app where physical appearance is a key part of how users understand each other. But even if users end up sharing information that almost exactly mimics the existing profile, we'll be able to conclude that those characteristics, terms, and labels are in fact representative of how gay men understand their bodies and identities. A free-form profile design has the additional benefit of allowing profiles to reflect culturally diverse understandings of how bodies and identities can be expressed online; culturally specific details that don't currently have a place in a Grindr profile might become a key part of how the service is used across different locations. The overall benefit is increased freedom of expression.

There are two significant tradeoffs to this approach. First, a free-form profile creation process potentially imposes significantly more creative labor on users themselves. The current Grindr profile is built to minimize friction: that is, to make the profile creation process as quick and easy as possible, while gathering all the relevant information from users. Users merely need to answer the questions presented to them in order to create their presence on Grindr. A free-form process, by contrast, requires each user to consider and type into being the details of his profile — a more time- and effort-intensive task of creative labor that doesn't necessarily provide significant rewards. Second, these changes potentially undermine one of key affordances of the Grindr

platform: the ability to quickly filter visible users on the basis of the characteristics expressed in profiles. The sophisticated search and filtering tools available in the app rely on different users' data taking roughly the same form; eliminating the structure of the profile creation process makes filters considerably more challenging to implement.

In the end, a more conservative approach to diminishing structure in the profile creation process might be warranted. In place of pre-defined body types and Tribes, text entry fields could allow users to craft their own definitions of self — or, perhaps, to collaboratively author new signifiers through shared use. These changes are part of a balancing act between freedom of expression and ease of use that requires a tremendous amount of care on the part of interface designers. But the fact that these changes were successfully implemented at the scale of Facebook's billion-plus user base suggests that other service providers should not shy away from asking whether their interfaces can be crafted to better reflect the diversity of their users. The designers responsible for profile creation processes should continually experiment with new options for user expression — striving, continually, for a perhaps unattainable perfect fit.

CHAPTER 3

LIFE

Manhunt is one of the most popular members of what I have termed the second generation of gay-targeted networked media, boasting over six million active members. Its design and practices, alongside services like Gaydar and Gay.com, helped shape the direction of contemporary platforms like Grindr. One of these practices — the new user registration process — offers us an important insight into what happens after the last chapter’s discussion leaves off. As part of the registration process on Manhunt, new users are encouraged to upload a photo of themselves for inclusion on their profile. But not any photo will do. Photos uploaded to Manhunt are not immediately made publicly visible; instead, they are first reviewed by a team of screeners (many of whom, a former employee told me, are heterosexual men and women) against a published set of guidelines for acceptable content in user photos. Photos can be classified as “green” (unambiguously in compliance with photo rules and acceptable for public display), “yellow” (in compliance with the rules for photos that are not publicly visible), or “red” (unacceptable for display). The whole process typically takes half an hour, after which acceptable photos appear on the user’s profile. The procedure is, from the user’s perspective, effortless.

But what processes underlie the user experience of uploading a new photo? This chapter examines the content management policies and practices of Grindr, in order to identify the technological, legal, institutional, and social affordances which enable the exclusion of certain types of images and behaviors from the service. What are the rules that govern the types of content that are permitted to be displayed on gay social

networking services? How do the policies implemented on Grindr compare to those on other gay-targeted geosocial networking services (such as Scruff), or those on second-generation, browser-based services like Manhunt?¹³ How does the fact that these services are accessible primarily through mobile applications on smartphones impact their content policies? Are these policies the product of institutional restrictions on the kind of content that can appear in smartphone “app stores,” or do they represent a particular vision of what gay social networking should look like? How are those rules deployed in practice to manage user behavior? What is the responsibility of service providers — to the law, to their users, and, perhaps, to gay communities at large? More generally, what are the sociotechnical norms that structure how online services function, and how they conceptualize their relationships with their users?

This relationship between individual users and online service providers is notoriously difficult to understand. Joseph Turow, writing about Facebook, has referred to online service providers as a “black hole” (Turow, 2012, p. 138), absorbing and acting upon user information without providing any real indication of the logic behind or implementation of their procedures. Service providers are often reluctant to divulge proprietary or potentially competitively sensitive information. In lieu of looking inside the black hole, therefore, this chapter describes the documents, policies, and practices at its periphery; it interrogates the limited information that service providers make available in order to draw inferences about the internal logic of their content management

¹³ While I primarily consider the policies and practices of Grindr, this section also includes a comparison to two other mainstream gay social networking services: Scruff (another geosocial networking app) and Manhunt (a primarily browser-based gay social networking service which has released a variety of different mobile clients). I use these points of comparison to differentiate between patterns of content management which are prevalent across a variety of services, as opposed to those which constitute specific interventions on the part of Grindr’s developers.

practices. This focuses on three specific questions: First, what are the principles of governance encoded in the text of terms of service documents and content guidelines? Using a legally-informed close reading of the text of these policies, this chapter outlines the principles of governance they encode. While imperfect and often obtuse, these texts offer a concrete articulation (though not necessarily the only articulation) of how service providers understand their relationships with their users. Second, what are the on-the-ground practices of users and service providers that operationalize the codified rules encoded in terms of service documents? Using both interviews with members of the senior staff of these services, as well as published commentary on their actions, this analysis contextualizes the formal policies, focusing on how real-world engagements with or applications of these policies correspond with or differ from their textual articulations. Finally, what are the popular discourses and practices that emerge around these policies? Drawing on mainstream print and online publications, as well as publicly-available user reviews and discussions of the three applications examined in this study, I outline the most prevalent discourses around these policies, focusing on how the relationship between policy and user behavior is described in media discourses about these services.

Managing user-generated content — and, particularly, photos — is by no means a concern unique to gay-targeted social networking services. Social networking services, including ones focused on romantic or intimate relationships, cut across gay and straight communities; but the policies in place to manage gay services are distinctive in their specificity. While other scholarship has examined the social practices that emerge around these services (Crooks, 2013; Gudelunas, 2012b; Mowlabocus, 2010a), few have engaged

directly with the distinctive normative and technical designs of these platforms. In evaluating the relationship between policy and practice on Grindr, this study outlines both a model of content policies at their most specific, as well as a model for how the relationship between technical systems and subcultural practice should be conceptualized. Framing content management policies as solely technical in origin obscures the value judgements that are embedded in them.

This chapter examines how the restrictive policies in place on services like Grindr came to be authorized, both by application developers and the users of gay-targeted social networking services. I argue that several factors are at play in these policies: first, ecosystem-wide restrictions on content that can be included in apps distributed through mainstream app stores operated by Apple and Google; and second, a set of clear normative interventions on the part of service providers into what gay social media ought to look like. While both factors contribute to the ultimate outcome of restrictive content management policies, I suggest that public discourse (cultivated in part by service providers themselves) overwhelmingly prioritizes the role of Apple and Google in restricting user content, while diminishing the role played by software designers and developers. The end result of these policies is a systematic constraint on the visibility of overtly sexual or non-normative forms of queer identity and self-expression through the promotion of a banal public face to gay networked media. I conclude by considering what I term practices of “expressive resistance”: that is, patterns of user behavior which push the limits of allowable content on gay social media through encoded communication. I suggest that these tactical engagements with content management represent the best

possibility for promoting free expression within the broader normative framework of mainstream gay platforms.

Objectionable, indecent, and pornographic

In February 2012, a former employee of oDesk, a company contracted to screen user-generated content on Facebook, leaked a copy of the service's operating guidelines (A. Chen, 2012). Content screeners are instructed to review any material that Facebook users have flagged as offensive or inappropriate and determine whether an actual violation of Facebook's Community Standards (Facebook, n.d.) has taken place. The list of forbidden content is extensive, with 50 separate items for review in nine different categories, including "Sex and Nudity," "Illegal Drug Use," and "Graphic Content." Among the myriad types of banned content are depictions of "any obvious sexual activity" (including in instances where no nudity is actually displayed), "female nipple bulges," urine, feces, vomit, semen, pus, and earwax. The guidelines also ban photos of mothers breastfeeding, a policy that prompted a widespread backlash against the service in 2008 (Calhoun, 2008; Ibrahim, 2010). The real absurdity of the guidelines, comments Tarleton Gillespie (2012), is the need "to draw this many lines in this much sand."

The leaked guidelines are compelling because they offer a rare look into the otherwise closed system of content review on Facebook. The process by which content is evaluated is completely opaque to users: The Community Standards — the only guidelines offered to users about what they may or may not post — are notoriously vague, and content disappears from view without explanation when moderated. The

reasoning behind these content restrictions is likewise seldom made clear to users, though compliance with state and national laws is high on the site's list of priorities.

Whatever the reasons for the restrictions, the fact that Facebook *at all* limits the content users are permitted to post has the effect of constituting a normative intervention on the part of the service into broad cultural controversies.

When Facebook steps into these controversial issues, decides to authorize itself as custodian of content that some of its users find egregious, establishes both general guidelines and precise instructions for removing that content, and then does so, it is not merely responding to cultural pressures, it is intervening in them, reifying the very distinctions it applies. (Gillespie, 2012)

This isn't to advocate an "anything goes" policy; limiting certain types of patently reprehensible content (child pornography, rape, bestiality, and self-mutilation are classic examples) is not only legally required but also generally ethically unproblematic. But reasonable or not — ethically, legally, and technologically justifiable or not — the fact that Facebook declares certain types of content off-limits at all is normatively motivated, whether or not Facebook's developers and users recognize it as such. The problem, suggests Gillespie, is that those norms are never made transparent to Facebook's users. In the classic words of Lawrence Lessig (1999), code is law — and content management is represented to its users as a set of encoded, technical rules. But Lessig's account, focused primarily on political structures relating to the operation of offline law in networked spaces, doesn't go far enough. As Wendy Chun puts it, in practice, code is better than law:

[Code is] an inhumanly perfect "performative" uttered by no one. Unlike any other law or performative utterance, code almost always does what it says because it needs no human acknowledgement ... Moreover, whereas a

law's effectiveness depends on enforcement (self- or otherwise), code's enforcement stems from itself. (Chun, 2006, p. 66).

By eliminating the speaker (the developer or designer of an application) from the equation, code reifies rules of conduct in a way that forecloses most opportunities for disagreement. We forget that those rules of conduct — the principles encoded in software through programming — had their origin in the mind of a developer, and accordingly represent a particular ideologically-motivated vision of how an application or service should operate.

The content policies of the gay-targeted social networking services examined in this chapter are a good deal narrower in focus. Yet, as is the case with Facebook, embedded in each of their policies is a core set of frequently recurring normative positions governing what content is considered “acceptable” within the context of gay-targeted social media. Excavating those positions requires a closer examination of the guidelines each service makes available to its users.

Interpreting these policies requires an understanding of their various roles in governing the relationship between services and users. Content management policies rest at the nexus of two sets of standards: first, that which is lawful; and second, that which is *proper*, as determined outside of and beyond the law. Many of the published guidelines focus on ensuring the legal status of these services, such as prohibitions on mentions of recreational drugs or escorting and “massage” (typically code for solicitations for sex for money). But in each case, the TOS documents for Manhunt, Grindr, and Scruff establish a broader class of content that is forbidden *within the context of these services*, regardless of legality. This is not just a question of legal rhetoric; it is a normative declaration that

what the law considers objectionable is, in some instances, not sufficient to govern online services.

Each service begins with an extremely broad set of proscriptions on user behavior. The Grindr terms of service offer a representative example: Grindr users are prohibited to

post, store, send, transmit, or disseminate any information or material which a reasonable person could deem to be objectionable, defamatory, libelous, offensive, obscene, indecent, pornographic, harassing, threatening, embarrassing, distressing, vulgar, hateful, racially or ethnically or otherwise offensive to any group or individual, intentionally misleading, false, or otherwise inappropriate, regardless of whether this material or its dissemination is unlawful. (Grindr, 2014a)

Similar language appears in both the Scruff “Terms of Service” and the Manhunt “Terms of Access and Use” (Manhunt, 2009c; Scruff, 2012). These documents establish a category of prohibited content which includes material which is “objectionable,” “indecent,” “pornographic,” “embarrassing,” or “otherwise inappropriate, regardless of whether this material or its dissemination is unlawful” — a vast category of potential content whose terms are often left ill-defined. To take up only one example, while indecency has been adjudicated in a general sense over several decades of American jurisprudence (Shafer & Adams, 2005), its specific meaning in the context of these services is left unstated.

What do “vulgar,” “obscene,” “objectionable,” and “indecent” look like in practice on these services? Manhunt, Grindr, and Scruff each translate their legal TOS and EULA documents into a set of operational guidelines for end users, often in significant and graphic detail (Grindr, n.d.; Manhunt, 2009a; 2009b; Scruff, n.d.).

Manhunt, for example, delimits four categories of photos that are prohibited on the site:

photos depicting scatology, urination on a person, blood, or weapons. Photos depicting “mid-stream urine or urination on inanimate objects,” as well as anal insertion and semen “in or around [an] orifice” are allowable, but only if designated as “private” photos, hidden from the general public of Manhunt.

The photo guidelines for Grindr and Scruff are considerably more restrictive. Scruff prohibits all instances of below-the-waist nudity (or partial nudity), including any exposed pubic hair or the display of genitals that are “obscured with hands, towels, hats, or by other means.” Visible erections or “tenting” in one’s clothing that suggests the presence of an erection is likewise forbidden. Grindr adopts a more general approach, writing simply, “No sexually explicit, revealing, or overly suggestive photos of any kind.” The guidelines then elaborate that skin below the hip bones, exposed underwear, and sexually suggestive objects (Grindr’s now-notorious “fruits and veggies” rule) are not allowed to be displayed in photos. Grindr also prohibits disclosing the size of one’s genitals in the text of a profile, as well as any other references to sexual acts.

The mechanics of content management likewise differ from service to service. All three services encourage their users to police each other and report photos or profiles that violate community guidelines — a common practice across social media services (Albrechtslund, 2008). Manhunt and Grindr, however, go further. Manhunt, for example, requires that every uploaded photo be screened before it becomes available on the site. The service’s published guidelines make careful note of the fact that the list of restrictions available to users is not exhaustive. The final determination of whether photos are “green,” “yellow,” or “red” is made at the discretion of the Manhunt staff during the photo review process. Grindr, like Manhunt, screens every uploaded photo

manually, using a team of photo reviewers (“censors,” as Grindr’s CEO described them in an interview; (Easton, 2009)) to determine whether a photo violates the service’s profile guidelines. Users are not given access to the internal logic of the review process (including in the event that a decision is made to reject their content), nor do any of the services publicly provide information about who the reviewers responsible for these determinations are. A former Manhunt employee noted, anecdotally, that many of the service’s contracted photo screeners are heterosexual males; but, beyond anecdote, none of the three services were willing to disclose further demographic information about the people in the screening role. Users are encouraged to focus on the outcome of the screening procedure, rather than the process itself.

Some of this variability can be accounted for by considering the technical context of these services. Unlike multi-device social networking services like Facebook, Grindr can be accessed exclusively through mobile applications, distributed through application distribution platforms like Apple’s App Store or Google Play. Accordingly, mobile-only or mobile-first social networks like Grindr




Fig. 5: Placeholder displayed while user profile content is under review on Grindr.

are bound by rules set forth by these distributors. Grindr’s developers emphasize the fact that their content guidelines are designed primarily to ensure compliance with the rules set forth by Apple and Google for developers on their respective mobile platforms. In discussing his app’s restrictions, for example, Grindr’s founder and chief executive Joel Simkhai begins by noting that, “Apple does not allow any nudity or profanity” (quoted in Easton, 2009), positioning Grindr’s specific policies as derived exclusively from externally-imposed restrictions — even as Apple’s restrictions (discussed below) do not proscribe nudity or profanity in all instances.

The rhetorical logic herein is clear: Developers remind their users that apps distributed through mainstream smartphone application distribution platforms require more restrictive content standards. The narrative presented to users is that, faced with the choice between not offering an application at all or abiding by Apple and Google’s rules, developers have opted to limit the types of content available on their services *for the users’ benefit*. This focus on externally imposed developer guidelines constitutes an important reframing of the discourse around content management practices. In particular, it shifts the responsibility for these policies off of application developers and onto Apple and Google. As Grindr’s Joel Simkhai explains, “From day one, we basically used the App Store guidelines as a framework for development.” The ambiguity of these guidelines, Simkhai continued, explains the Grindr staff’s cautious development approach:

Apple and Google don’t have very specific guidelines — sometimes they can be quite vague. Trying to make sense of them is often a Talmudic exercise, so when we drew up the Grindr profile guidelines, we were very conservative in our interpretation of Apple and Google’s guidelines. (Y. Roth, personal communication, December 9, 2012)

By focusing on the rules set forth by Apple and Google, Simkhai downplays the internal design process behind Grindr as a factor in developing content restrictions. The choices were made *for Grindr* by Apple and Google, rather than by Grindr's staff for their users. This explanation has caught on in popular discourse about the restrictions. For example, writing about an updated version of Grindr, the blog Queerty notes,

Software makers revise their guidelines all the time, but nobody tightens the rules faster than developers subject to Apple's increasingly stringent rules about what can be sold in its iPhone app store. ... [Grindr] has, because of the App Store's existing rules, never allowed members to display naked photos in their main profile pictures. But updated rules go much farther. (Queerty, 2010)

In 2010, when Grindr's amended profile guidelines went public, Apple had not substantially modified any section of its published developer guidelines, including sections that would affect services like Grindr. Most observers of Grindr's policy change presumed that Apple was responsible for the tightened restrictions, and Grindr's developers did little to dispel that assumption.

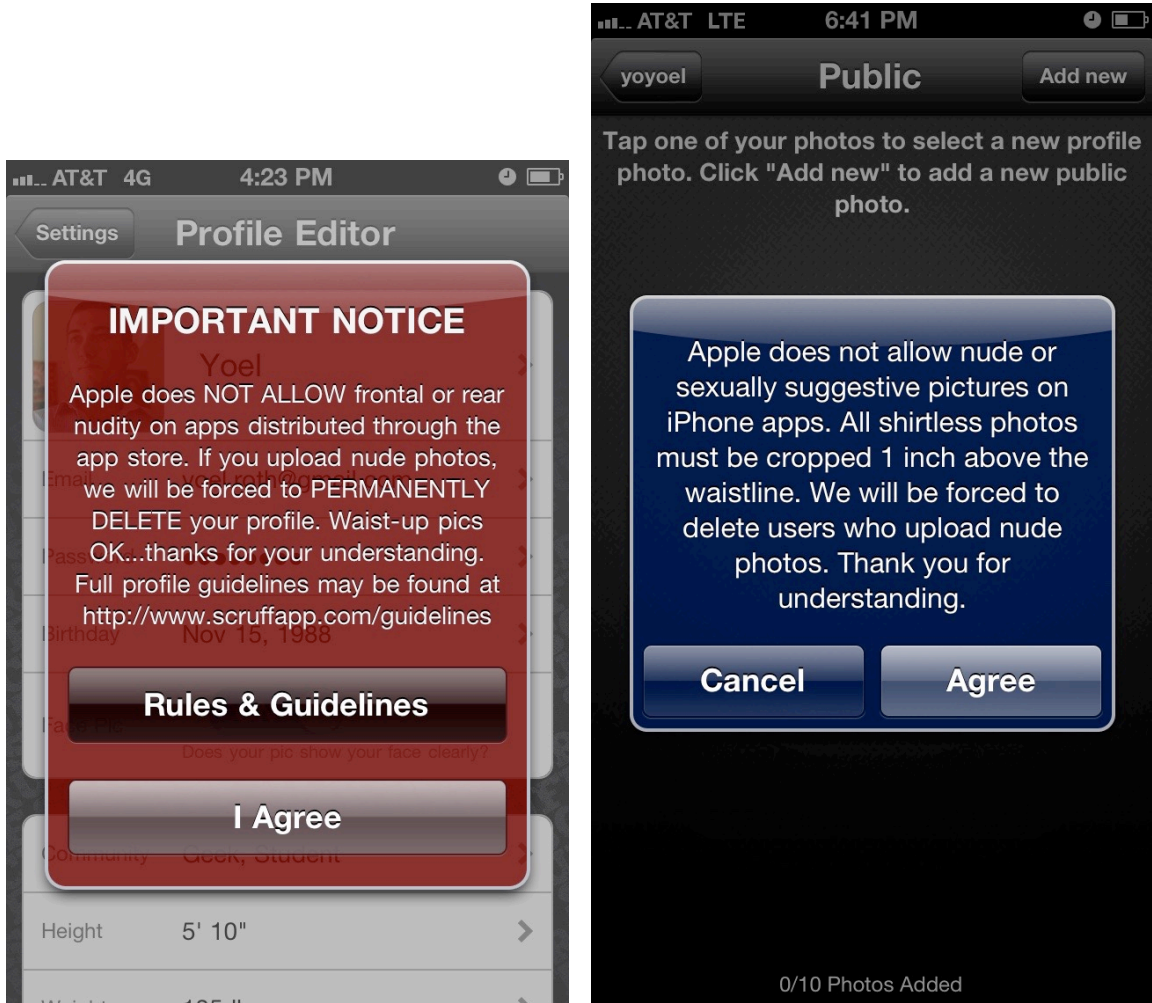


Fig. 6: Content policy notice on Scruff. The screenshot on the left, taken in August 2012, displays the notice used from the launch of the app. The screenshot on the right, taken in November 2012, uses revised language.

Scruff likewise places the responsibility for its content management policies on Apple’s developer guidelines. For example, the app displays a prominent message whenever users start the process of uploading a new profile photo, indicating that Apple prohibits certain types of content in user-submitted images, and that Scruff is required to enforce those prohibitions with reference to user profiles. The language of this message has evolved over time: from a ban only on frontal and rear nudity, to a blanket prohibition on “nude or sexually suggestive pictures.” Yet, as was the case with Grindr,

it's unclear what prompted Scruff to revise the text of its content policy notice in 2012; at least publicly, neither Apple nor Google had tightened the rules around sexual content in apps.

But what do Apple and Google's developer guidelines actually say? Are they sufficient in themselves to account for the policies in question? Apple's developer guidelines are a combination of practical injunctions and broad ideological statements:

We view Apps different [sic] than books or songs, which we do not curate. If you want to criticize a religion, write a book. If you want to describe sex, write a book or a song, or create a medical App. It can get complicated, but we have decided not to allow certain kinds of content in the App Store. ...

We will reject Apps for any content or behavior that we believe is over the line. What line, you ask? Well, as a Supreme Court Justice once said, "I'll know it when I see it". And we think you will also know it when you cross it. (Apple, 2013a)

The guidelines are more specific on the issue of pornography, noting that apps containing objectionable, crude, or patently pornographic material (user-generated or not) will not be distributed through the App Store. Google likewise notes that pornography, nudity, graphic sex acts, and sexually explicit material are all prohibited in applications distributed on Google Play (Google, n.d.). While potentially ambiguous, these policies do not, in themselves, prohibit the full spectrum of content addressed in the Grindr guidelines.

We can account for the gap between platform policies and specific app practices in two ways. First, as Luis Hestres (2013) argues, these policies constitute an important and non-content-neutral restriction on developer behavior. This is the most direct explanation of control: Apps constrain user behavior because Apple and Google

specifically proscribe certain types of content. But, in practice, these restrictions tend to be considerably more restrictive than direct control can account for. Instead, I would suggest that these platform-wide policies can create a chilling effect on developer behavior: Rather than running the risk of violating platform rules, developers elect to be more conservative in their specific policies. This is the explanation offered by Grindr's staff. Few services, however, acknowledge their own normative interventions into this process: An important act of translation occurs between the Apple and Google developer policies and the rules users actually engage with. Herein, both platform curators like Apple and Google *and* application developers behave in a non-content-neutral manner. And, in practice, this results in content policies that are more specific or comprehensive than a hypothetical least-restrictive-alternative that would comply with the strictures of platform rules.

Critically, we should take note of the fact that using Apple's developer guidelines as the justification for restrictive content policies is not a practice limited to the domain of gay-targeted social networking services. Speaking at an event in September 2015, Instagram CEO Kevin Systrom noted that controversial content management decisions on the part of Instagram's administrators were directly the result of the service's desire to comply with Apple's developer policies (Slater-Robins, 2015). Instagram, notably, has been the target of the #FreeTheNipple movement, in which users argue against a seeming double standard wherein male nipples are not censored on the platform, but female nipples are. In discussing the Facebook-owned platform's policies, Systrom sidestepped the particulars of #FreeTheNipple and argued more generally that keeping the Instagram

app rated “12+” in the App Store requires some concessions; in his words, “in order to scale effectively, there are [some] tough calls” (Slater-Robins, 2015).

The opacity of the app review process leaves little room for outside observers to contest Systrom’s account of the provenance of Instagram’s policies. Perhaps, as Systrom suggests, Instagram was threatened with removal from the App Store if it failed to sufficiently protect under-18 users from a barrage of sexually-suggestive images. But, in both the cases of Instagram and Grindr, we should question why these apps in particular seem to have run afoul of the Apple App Store review process, whereas others have not. Notably, the Twitter app has maintained a “4+” parental rating in the Apple App Store, despite the facts that users under 13 are not permitted on the platform at all, *and* that Twitter is host to a substantial quantity of readily-accessible sexual content (Kleeman, 2015). Why are Instagram and Grindr seemingly held up to stricter scrutiny than Twitter? One possibility, as Hestres (2013) suggests, is that the Apple app review process is incurably arbitrary; perhaps Grindr and Instagram have been unfairly treated differently from their peers on the App Store. Another possibility, and one which in the remainder of this chapter I suggest is the case, is that Grindr and Instagram have voluntarily adopted non-neutral approaches to user-generated content which constitute active normative interventions. I argue that these policies, whatever their basis in app store developer guidelines, are reflections of closely-held but seldom-revealed philosophies about the ideal characteristics of a social networking service.

Normative platforms

In the early days of Web 2.0, interactive electronic content was seen as a boon for the agency of individual users. Free expression — of identities, of diverse viewpoints, of artistic creations — could be facilitated through open access to online services. As Andrew Barry optimistically put it, the logic of online interactivity is “You may!” not “You must!” (Barry, 2001, p. 149). But in practice, “may” and “must” have turned out to be two sides of the same coin. User behavior online has not only tended to be constrained, but in fact is engineered from the ground up in a manner that cannot help but be constrained (Gillespie, 2007; Lessig, 1999). Those patterns of constraint emerge in the content restrictions examined in this chapter. The key to understanding these restrictions is to position them in the broader sociotechnical context that authorized their creation in the first place.

Tensions over the normative characteristics of electronic services have tended to cluster around the management of user-generated content. The politics of content management play out in the discursive construction of platforms as a way to describe online service providers. The general characteristic of a platform is that it is an “open, neutral, egalitarian and progressive support for activity” (Gillespie, 2010, p. 352) — that it exhibits the quality of content-neutrality while giving users a medium for self-expression. An ideal platform, to return to Barry’s phrase, tells users, “You may!” and gives them the tools with which to do so. In practice, however, the term “platform” has been applied to services that only rarely exhibit content neutrality and openness. In José van Dijck’s words, oftentimes a platform “shapes the performance of social acts instead of merely facilitating them” (van Dijck, 2009; 2013b, p. 29).

The idea of a platform has its own particular set of affordances. In particular, calling a service a “platform” downplays the responsibility of a service’s owners and developers for the content they distribute:

Online content providers who do not produce their own information have long sought to enjoy limited liability for that information In the effort to limit their liability not only to these legal charges but also more broadly to the cultural charges of being puerile, frivolous, debased, etc., intermediaries like YouTube need to position themselves as just hosting — empowering all by choosing none. (Gillespie, 2010, p. 357)

By referring to a service as a “platform,” the burden for choices about content restrictions is shifted from the developer onto larger structures of constraint, like legal requirements and nebulous social standards like “common decency.” More generally, the rhetoric of platforms makes content management practices susceptible to what Zygmunt Bauman (1995) has termed “adiaphorization”: that certain activities, particularly by organizations, are rendered as outside of the sphere of moral or ethical evaluation. The rightness or wrongness of the rules in question becomes irrelevant; instead, “It’s not my department” becomes a characteristic means for deferring, potentially indefinitely, responsibility for the normative outcomes of enforced policies (Bauman & Lyon, 2013, p. 13; Clegg & Rhodes, 2006, p. 7).

I argue that these practices of normative deferral become yet more problematic when they take place in a digitally mediated institutional context. The image approval processes of services like Manhunt and Grindr are presented as constraints resulting from technical systems even as they manage user-generated content by encoding human norms about controversial or sensitive material in only quasi-technical terms. But this humanness is seldom revealed to users. The reviewers responsible for screening

submitted photos operate in a black box. Photos, as if by magic, are approved, rejected, classified, and edited, with no revelation of the underlying mechanism. Hints of the human underbelly of technical platforms are revealed only by accident, as in the case of the leak of the Facebook content guidelines.

At stake in these conversations is the place of values in the design of online services. As Helen Nissenbaum valuably recognizes, accounting for the operation of values in and through technology requires a broad analysis of the interplay “between the system or device, those who built it, what they had in mind, its conditions of use, and the natural, cultural, social, and political context in which it is embedded” (Nissenbaum, 2001, p. 120). Many of these factors are challenging to account for — Can we ever accurately discern what the developers of a technology truly had in mind? — but acknowledging their presence in the design process denaturalizes the choices that contribute to the emergence of certain types of technological systems. The developers of technologies have a corresponding ethical obligation to explicitly integrate some consideration of social values into their design processes (Flanagan et al., 2008; Friedman, 1996; Friedman & Nissenbaum, 1996).

Conceiving of social services as value-neutral platforms has the effect of creating a “comforting sense of technical neutrality” (Gillespie, 2010, p. 360). But this obscures the broader political questions at stake: Do these services have *responsibilities* — and if so, to whom? The desires of a majority of their users? The law? Their shareholders? Where, if anywhere, can the particularities of political conflict, history, and subcultural construction be integrated into this understanding of how online service providers operate? Striking this balance is an inherently non-neutral process. Reconceptualizing

these choices as the active negotiation of social responsibility would enable service providers to integrate value-sensitivity into their design process in a more robust way than the doctrine of platform neutrality allows.

Content management policies are an important instantiation of non-neutral design choices. Where the normative character of these policies is made properly visible to users and developers alike, it affords them the opportunity to engage with those struggles in a robust, value-centric way. The neutral rhetoric of platforms, by contrast, privileges the technological status of these services over their cultural valences. Value-sensitive design is of particular importance as an analytic framework when traditionally marginalized values are more contextually prominent than others, as is the case on these services.

Negotiating gay visibility

In the cases of Grindr, Scruff, and Manhunt, a particular set of histories and struggles over the visibility of queer sexualities are implicated by virtue of their status as *gay-targeted* services. Gay-targeted online networks have existed for more than two decades, beginning with social spaces like Usenet groups (O'Riordan, 2005) and Internet Relay Chat rooms (Campbell, 2004). The services examined in this study constitute the second and third generations of gay social media — services which are built around allowing users to express themselves with higher-bandwidth media like photos (as opposed to only text), and which take advantage of emergent technologies like geolocation and always-on mobile data connections. Yet, as gay-targeted social media have matured technologically, they have also become more restrictive of the types of content their users are permitted to

share. This normative question has significant historical roots in and ramifications for Western gay politics.

The practices of content management taking place on gay-targeted social networking services parallel long-standing offline debates over the visibility of gay sexuality. Tearooms, bathhouses, and cruising areas — sites where certain types of gay sexuality can become visible to the public — played a significant role in structuring gay politics throughout the twentieth century. The disavowal of highly visible forms of gay sexuality by mainstream gay organizations represented a significant and controversial turning point in the history of American gay politics. A wide range of scholarship has chronicled various facets of this shift — from public health to urban zoning laws to the emergence of a new “homonormativity” that privileges private, moderate, and often implicitly seronegative homosexuality (Berlant & Warner, 1998; Delany, 2001; Elovitz & Edwards, 1996; Seidman, 1992; Warner, 1999).

A related perspective, discussed in Elija Cassidy’s research into gay men’s use of Facebook and Gaydar, suggests that these tendencies towards moderation and the careful sequestration of overt sexuality have resulted in a default of banal homosexuality on social platforms. Cassidy writes that, on Facebook,

information about what [someone] ate for lunch, pictures of a newborn family member, how much they wished they could go home from work, and where they wanted to go on their next holiday, was simply presented within the framework of a profile which, for example, linked them to a male partner, a gay male interest group, or listed a gay icon as their “Religious Views.” (Cassidy, 2013, p. 120)

Cassidy contrasts the banality of sexuality on Facebook with the more overt and sexually-explicit profiles his informants maintained on Gaydar. Different platforms, Cassidy

suggests, promoted different forms of disclosure: some manifestly sexual; others more quotidian or oblique. Individuals take their cues about how to share information about their sexuality from the affordances of a given platform, limiting their disclosures to those which correspond with the expected conduct of the service in question.

Grindr and Scruff represent especially interesting sites for questioning the banality of sexualized self-expression on gay networked media. Whereas Manhunt's more direct self-positioning as a venue for finding casual sex locates the service within the tradition of platforms like Gaydar and Adam4Adam, Grindr and Scruff maintain a consistently agnostic stance towards the activities of their users. A wide range of studies have indicated that the motivations of the users of gay-targeted social networking services are often diverse and not restricted to the solicitation of sexual partners, though "hooking up" is an important and frequently-cited use (Crooks, 2013; Gudelunas, 2012b; O'Bryan, 2012; Vernon, 2010; Wortham, 2013). Grindr's popularity is at least in part a product of enabling that diversity by embracing the polysemy of the service: Grindr's developers refer to the service as a "gay friend finder," cautious and possibly euphemistic rhetoric that leaves the actual negotiation of the service's use open to users. Users are left to determine for themselves the character of their in-app interactions.

Content management has been deployed strategically to contribute to the ambivalent definitions of purpose around Grindr and Scruff. Profiles are permitted to be suggestive — but only to a point. As the *New York Times* put it in a profile of Grindr's Joel Simkhai,

"I see us more as a bar than a sex club," Mr. Simkhai said. "If you go out to a bar, you don't want to see someone with his genitals hanging out." And if, as Mr. Simkhai said, there would be a "certain ickiness" to Grindr

devolving into a mere digital sex club, that is not to suggest the desired endpoint for him or any other user is to organize a holiday food drive or a Scrabble tournament. (Trebay, 2014)

Herein, Simkhai makes his normative position with regards to sexually explicit content clear: such images would be disruptive to the open-ended social space he imagines Grindr to be (though, of course, following from his statements there's no reason a user *couldn't* use the app to organize a food drive). Grindr's content management strategy attempts to establish a point of equilibrium between the "ickiness" of a digital sex club and the prudishness of a Scrabble tournament. The underlying assumption is that overt displays of sexuality would alienate users who *aren't* on the hunt for hookups. Or, in Simkhai's words,

We didn't want our users to experience logging into Grindr and instantly seeing unexpected or unwanted nudity. ... We want Grindr to be inclusive. A lot of the services that came before Grindr were more overt, and we think that narrowed the appeal and the experience of using those services. (Y. Roth, personal communication, December 9, 2012)

The key to making sense of these choices is to recognize that they represent an attempt to define, in a general, qualitative sense, what kind of space a given application is supposed to be. Simkhai suggests that part of Grindr's viability as a social network is based on keeping erotic images hidden from view. Content that falls outside of the boundaries of the Grindr profile guidelines is taken at face value by Grindr's developers to be disruptive to the safe and enjoyable operation of the Grindr service.

This isn't to say that every social networking service is obligated to subscribe to a version of gay politics that prioritizes highly visible sexuality. Gay social networking services are not a one-size-fits-all proposition. This logic is not new, or specific to gay social networking services (boyd, 2007). Different social networking services may appeal

differently to particular segments of a population. Users who don't fit on omnibus services like Grindr can jump ship to their competitors, like Scruff or Mister. Larger-bodied or more hirsute men may be drawn to Scruff instead of Grindr (Roth, 2014). Older men might prefer Mister. Ultimately, users of gay social networking services pragmatically seek out the applications that “work” for them — ones that offer the greatest number of social gratifications with the fewest burdens or barriers to participation.

None of these services, including Grindr, are developing software with the intention of marginalizing certain types of people. Inclusivity of a wide range of users, ease of access, and ubiquity unsurprisingly emerge as the dominant discourses in the marketing materials of each of these services. As Manhunt's slogan puts it, “If he's out there, he's on here” — implying that users will have a better experience if the network they're a part of includes as large a segment of the gay population as possible. This logic directly informs these services' approach to content management. At the core of their practices is the belief that the pursuit of a broad base of potential users and the overt display of sexual (or even highly suggestive) content are mutually exclusive.

Whatever the actual use-cases of these apps, users have rejected the premise that effective gay social networking requires hiding sexual or suggestive content from public view. Many cite these policies as disruptive to their experiences using these apps. In the case of Grindr, for example, the service's content policies are a frequently-recurring theme in the more than 30,000 reviews of the application in the Apple App Store. Many users refer to content policies as a form of censorship on the part of the app's developers. These policies, users contend, are at odds with the idea of a social service targeted at

adults. “I can understand no nudity,” writes one user, “But these guys go overboard.” The seeming arbitrariness of the guidelines — and the fact that other social services do not similarly restrict user content in such a granular manner — appeared often in users’ complaints. As one user notes, “If Facebook approves a photo ... shouldn’t a site like this?” These accounts have been repeated in mainstream and online press coverage. An article published in *Vanity Fair* referred to Grindr’s content policies as akin to “the student handbook at a parochial school” (Kapp, 2011), and a number of the most trafficked gay blogs have called the practices “puritanical” or “prudish” (Easton, 2009; Queerty, 2010; Towle, 2010).

Despite public outcry about the policies, users have been ineffective in prompting widespread changes in policy. Developers are able to resist public pressure by framing the rules of their services as technically-derived, rather than normative (and therefore open to contestation). Users themselves occasionally repeat these explanations. As one reviewer of Grindr notes, “Most of the negative reviews here ... seem to be primarily by folks expecting some sort of x-rated free for all on an app store product.” In this account, the App Store, not Grindr, is responsible.

Additionally, the all-or-nothing approach of these policies precludes many forms of resistance by definitionally excluding users who disagree. Users can either accept the terms of service, or opt-out altogether. In the words of a Grindr user, “The worst part of the app is that it might be the best app like it in the app store.” Despite frequently-expressed and often quite vehement criticism, users decline to opt-out. Opting-out as a strategy of resistance has its own corresponding set of constraints. Peer pressure and collective inertia create significant disincentives to “vote with one’s feet” and leave a

particular service. The dominance of a small number of services means that electing to use their less popular competitors necessarily reduces the number of available connections. As another user put it in a review, “Why do we all keep using this app? Let’s face it: the only reason to use this app is because it has so many users.” The fact that millions of users continue to do business with Grindr should not be taken as a strong indicator of their approval of the service’s practices; at best, it demonstrates that the disincentives of leaving are not outweighed by the potential benefits of networking elsewhere.




In lieu of opting-out altogether, users often elect to use particular gay-targeted social networking services for contextually-specific purposes. David Gudelunas (2012b) reports that, among his respondents, some indicated having seven to ten active profiles across different services. These profiles are not redundant. Instead, users present themselves in different ways on different services, both in accordance with the rules of the service and based on their own expectations of what types of interactions a particular network has to offer. These profiles, Gudelunas notes, “were not seen as discrete entities, but rather as part of an elaborate network” — a fragmented, yet coherent, version of individual identity online that accounts for the differential permissions and gratifications afforded by social services. Opting-out altogether is a last resort.

Expressive resistance

In addition to using multiple services in combination to achieve different goals, I want to highlight a broader set of negotiations that take place within the context of a particular service. Specifically, I argue that users are not completely disempowered by

sociotechnical constraints — including specific constraints on the content they are permitted to share. Rather, their behavior is dialogically negotiated on an ongoing basis, often in ways that neither the user nor the developer would initially anticipate (Best & Tozer, 2013). Some users creatively engage with the rules of technologies they encounter in order to make externally imposed boundaries more personally comfortable or agreeable. They can, borrowing from Michel de Certeau (1988), work tactically within systemic constraints to achieve a desired outcome. I term these practices “expressive resistance.”

Emoji — a set of 722 small, elaborate pictograms built into the iOS and Android operating systems — play an important part in the processes of expressive resistance. These pictograms were originally developed in Japan in the early 2000s by the mobile phone provider DoCoMo as part of an attempt to make mobile messaging more appealing to a teenage demographic (Lebduska, 2014). By 2010, emoji had been incorporated into the Unicode Standard (a governing protocol for the handling and display of text on computers and mobile devices), and were widely available on mobile phones outside of Japan. But even as emoji became a part of the formal Unicode standard, their meanings were never fixed. While some are direct descendants of textual emoticons, many remain ambiguous, giving users the possibility to deploy them as they see fit — rather than in accordance with a set of prescribed definitions. Emoji, writes Lisa Lebduska (2014), by design “open themselves to re-appropriation, interpretation, and even misinterpretation.” Their meaning, she suggests, is often a matter of context and shared cultural meaning, rather than an externally-imposed set of significations. This semiotic openness plays a crucial part in enabling expressive resistance.

While Grindr’s manual photo review processes leave little room for circumventing guidelines about images, users nevertheless display prohibited content on their profiles in only slightly disguised ways. For example, users employ specific emoji to identify as either the penetrating (“active” or “top”) or penetrated (“passive” or “bottom”) partner in anal sex — explicit references to sexual activity that are prohibited under the content guidelines. Through the use of the emoji “”, “”, and “”, users maintain compliance with the text of the content policies, even as their profiles convey information that would otherwise not be permitted. Users also employ encoded communications to efficiently convey information about themselves that does not have an immediately obvious place in the design of profiles on Grindr. For example, some seropositive users of Grindr work around the lack of a formal interface element for disclosing seropositivity (a feature available on many second-generation gay social networks, including Manhunt) by including a bracketed “[+]” in their profile headline.

Using Stuart Hall’s (1973) framework of encoding and decoding, we can understand the use of emoji and textual symbols on Grindr as part of a complex, multistep, multi-agent system of communication. Operating within a domain of shared knowledge (what Hall terms “frameworks of knowledge”), Grindr users encode the meaning of their messages (for example, their preferred sexual position) in a manner that, superficially, complies with the terms of Grindr’s policies. In successful communication (“meaningful discourse,” as Hall terms it), other users are able to decode these messages and act accordingly. The communicative process herein is attenuated, and relies on a body of shared knowledge about the meaning of particular symbols that is not readily

available to new users or community outsiders. There is not, for example, a beginner's guide to communication on Grindr, popular-historical overviews like Jamie Woo's *Meet Grindr* (2013) notwithstanding. These uses of encoded communication techniques only became available to me through participation in the Grindr community; I encountered their use, and gradually puzzled out their meaning. Nevertheless, in many cases, the obviousness of the symbols in question, or the ability to easily ask another user for clarification (without running afoul of Grindr's terms of service), makes communication through expressive resistance a generally successful way of negotiating with the constraints imposed by top-down content management strategies.

We should not mistake the use of encoded signifiers for the presence of truly free communication. The very act of encoding information in emoji creates a barrier to access for new users, requiring them to learn the symbolic parlance of a particular service before they can meaningfully participate in its community. Further, there's little stopping service providers from learning the codes and proscribing their use as well. In some instances, developers have responded to these encoded practices by integrating specific mentions of user vernacular into content guidelines. For example, across the services examined in this study, users deploy terms like "partying" (a general reference to the use of recreational drugs), "skiing" (a specific reference to cocaine), and "420" (a specific reference to marijuana) to make their interests legible to other users while not running afoul of laws or the rules of a service (Race, 2015). Manhunt, in response, has specified the grammatical circumstances within which the term "party" may be used, in an effort to combat drug references in user profiles: "The term 'party' isn't allowed when used as a verb or adjective, but as a noun such as in 'sex party, 'dinner party' or 'group party' it is

allowed” (Manhunt, 2009b). These guidelines recognize the possibility that, in online spaces, users may attempt to subvert services’ limitations on their conduct by encoding discussions of their practices in ways that are legible only to other members of the community — to circumvent the rules, albeit covertly. The use of emoji to share sexual details would be even more straightforward to prevent. As part of the Unicode standard, all characters (including emoji) have particular alphanumeric values; automatically detecting the use of emoji to express overtly sexual content on profiles would be trivial, from a programming perspective. Grindr and Scruff could, like Manhunt, identify and prohibit coded linguistic tactics of user misbehavior in their terms of service.

Crucially, however, Grindr and Scruff have *not* banned the use of emoji to communicate sexual details in user profiles. The specificity of many of the content policies indicates that developers have invested a great deal of time and effort in understanding what users do on their services, and how best to govern their conduct. Developers are not ignorant of what their users are doing; instead, in select instances, they elect to behave as though they are. The result is what Foucault termed a “margin of tolerated illegality” (Foucault, 1978a, p. 82): a space within which institutions of power permit (through non-enforcement) their subjects to act in a manner which seems to contradict official rules. The fact that Grindr, Scruff, and Manhunt are vigilant about enforcing some content restrictions while turning a blind eye to others suggests that the normative prioritization of certain values is inseparable from the presumptively “neutral” practices of platform stewardship. Certain practices are selected as worthy of specific proscription in content guidelines, whereas others go unacknowledged — and therefore are permitted by default (until or unless a service’s developers specify otherwise). These

encoded communications offer users an important degree of freedom in the face of imposed constraints: The limitations of a particular service — whether created by explicit proscription in content guidelines or through the omission of a particular feature — nevertheless leave room for a limited amount of user negotiation. On the whole, user-generated content is still tightly controlled, and certain practices remain categorically marginalized; but the highly specific policies these services adopt are effective in creating a zone of plausible deniability within which user behavior may be relatively less constrained. The specificity of technical management gives rise to the possibilities of individual resistance.

Conclusions: The least restrictive alternative?

Services like Grindr are increasingly regarded as the new popular front of gay sociality. But what values are embedded in the networks created by these applications — and what are their consequences for gay communities, both on- and offline? Even as these services bring together millions of users from across the globe, certain elements of gay culture have been systematically pushed to the margins. Apps like Grindr make visible to the public a face for gay sociability that's devoid of the kinds of highly visible sexuality that were characteristic of Western gay culture in previous decades. Non-normative practices — fetishistic, “unsafe,” or highly visible sexualities, for instance — are consistently hidden from view on services like Grindr. Where non-normative or sexual discourses persist, they only manage to do so as a form of tactical, expressive resistance, visible only to those who speak the coded language of the app's community. It's perhaps more than just coy wording at work when app developers refer to these services as “friend finders.” By

consistently restricting the display of certain types of sexual content, social networking services like Grindr and Manhunt are entering into widespread debates over the banality of gay identity online — whether or not they acknowledge their actions as doing so. Content management is an essentially political process, and as the user base for these services continues to grow, the stakes for this process will only become greater.

The technical systems that frame these services obscure the relations of power encoded in them. Presenting a content policy as the product of technological requirements rather than normative ones reduces opportunities for user resistance and self-expression. Hiding certain types of sexuality from view on social networking services isn't to say that they don't exist; but diminishing their visibility is in itself a value judgement and an affordance for a particular and limited type of representation. These are normative considerations, not technological ones — considerations which have important consequences for the agency of individual users as well as the visibility of diverse practices and patterns of self-expression in gay communities.

As a matter of practice, however, it remains an open question as to how governance and content management should work on gay-targeted social networks in the future. One solution is to eschew app store distribution altogether, in favor of browser-based clients that do not have to comply with Apple and Google's content standards. In 2011, Manhunt implemented precisely this kind of change, offering users a mobile-optimized website as an alternative to its applications, complete with unrestricted access to all the user content available on the desktop version of the Manhunt service. In so doing, Manhunt opted out of the regimes of content restrictions that Apple and Google impose on developers who distribute their software through the App Store and Play.

While the browser-based Manhunt client is indeed less restrictive than the app store version, it's worth recognizing that the service's overall content restrictions and content management practices have remained unchanged. The rules Manhunt put into place to comply with Apple and Google's guidelines have not been loosened (Manhunt, 2011). A change in technical medium — from application to mobile web page — did not alter Manhunt's overall approach to user content. The service persists in establishing a hierarchy of user content in which overly fetishistic sex is forbidden outright, and activities that are perceived as potentially offensive or less safe are hidden from public view. Manhunt's developers are neither required by law nor by any external authority to limit what users are allowed to post, particularly along the highly specific lines they establish in the site's terms of service — yet they continue to do so.

Moreover, browser-based services (like Manhunt and a newer service called Squirt) have failed to attain the massive popularity and high levels of ongoing engagement that are characteristic of Grindr and Scruff. The shift from app to browser (in a sense, a reversal of the overall course of internet history) makes these services feel clumsier and harder to use, even as it affords the possibility for fewer content restrictions. Even when they have the option to use a less restrictive service, most people continue to use Grindr. This preference for ease of use and a preexisting large network should not be discounted.

The more obvious, and in some ways easier, response to the puzzling prudishness of these policies — and one which has been raised by countless users and scores of bloggers — is to implement less restrictive content management policies. It's difficult to justify a ban on photos of users in Speedos, or of users holding (suggestively or not)

fruits and vegetables on the basis of any available policy texts from Apple or Google; such images are not overtly pornographic, nor more revealing than one might expect from an album of family vacation photos. These images are banned on Grindr and Scruff because of their presumed context of overt sexuality, rather than because the content of the image is in itself indecent or objectionable by the standards outlined in the Apple developer guidelines. I would not hesitate to echo the calls from Grindr and Scruff users to revisit these policies and adjust them to become a least restrictive alternative: that is, a policy that keeps these apps in line with the rules set forth by Apple and Google, while permitting users the greatest amount of freedom to share information about themselves. This process of policy revision is complex and iterative, and is often bound up in inscrutable practices of review implemented by both Apple and Google for apps distributed through the App Store and Google Play. Nevertheless, I take it as a responsibility of the developers of gay-targeted software to make an active effort to craft policies which are as permissive as possible.

Policy revision is a delicate balancing act. I agree with Grindr's Joel Simkhai when he suggests that implementing policies based on a conservative interpretation of the Apple developer guidelines runs fewer risks of having the Grindr app unexpectedly removed from the App Store (thereby keeping scores of potential users from downloading it). And, certainly, the opacity and occasional arbitrariness of the app review process (Hestres, 2013) warrants a cautious approach. Ideally, one would hope to see a degree of flexibility — or even an openness to change — in Apple's handling of user-generated content in apps distributed through the App Store. A dialog between Apple and the developers who take advantage of its app distribution platform might result

in policies that better reflect the on-the-ground necessities of content management, rather than dogmatic rules that have remained relatively unchanged since the launch of the App Store in 2008. For developers familiar with Apple's management of the App Store, this scenario of transparency and an openness to negotiation seems unlikely. Cautious, iterative change represents the best available solution to the problem of content management. Yet, as a practical question, even subtle policy revisions seem unlikely. The disincentives to act — namely, the risk of getting booted from the App Store — are too significant to warrant changing policies that, thus far, have not significantly diminished the popularity of mainstream gay mobile apps.

Perhaps, in the end, the least restrictive possibility is the one which users themselves have put into practice: namely, encoded, expressive resistance that allows users to share the information that's important to them, without raising the ire of developers, censors, and policymakers. Expressive resistance is predicated on making do with the available tools of a given technical platform — limited or restrictive though they might be. It also gives developers the possibility to turn a blind eye to user activity, using plausible deniability as a middle ground between outright endorsement and active proscription. Users already live in this middle ground, and will continue to do so in the face of whatever future restrictions might be implemented on their behavior. To the greatest extent possible, therefore, software developers and policymakers should try to stay out of the way.

CHAPTER 4

AFTERLIFE

On October 1 2014, Dutch artist Dries Verhoeven logged onto Grindr. “Expect the unexpected!” he wrote in his profile. Verhoeven, a slim, bearded man in his late 30s, unsmiling and clad in a black t-shirt in his profile photo, seems like any other of Grindr’s millions of active users. But for many of the men he chatted with on Grindr between October 1 and 5, their interactions with the artist were anything but expected. At the time, Verhoeven was living in a glass-walled gallery space in the center of a public square in Kreuzberg, Berlin. He was connected to Grindr on five smartphones, and was projecting his interactions with other users onto the walls of the gallery space for passing members of the public to see. Verhoeven described the scene, titled “Wanna Play?”, as “an




Fig. 7: Public view of Dries Verhoeven’s “Wanna Play?” installation in Berlin, Germany, showing mostly unobscured personal information from Grindr profiles. Photo from Hebbel am Ufer, linked to at: <http://www.dazeddigital.com/artsandculture/article/22076/1/berlin-grindr-art-installation-shut-down-after-protests>

installation-performance that exposes the opportunities and tragedies of a phenomenon in gay culture: the sex date app.”¹⁴ The project, funded by the German avant-garde center Hebbel am Ufer, was designed to interrogate how emerging technologies like Grindr and Tinder “influence the way we present ourselves and connect to each other, both in a positive as [*sic*] negative sense” (Verhoeven, 2014). Verhoeven did not disclose to users — either in his profile or in private chats — that their interactions were part of a public art installation.

Three days into Verhoeven’s stay, he struck up a conversation with a local photographer named Parker Tilghman. Verhoeven invited Tilghman to come meet him. When he arrived at the address Verhoeven had given him, Tilghman found his conversations with the artist enlarged on the gallery’s walls. And, in his words (posted on Facebook), “I lost it” (quoted in Tharrett, 2014). Tilghman attacked Verhoeven, punching him and destroying furniture in the installation space. In the same statement, Tilghman described Verhoeven’s project as “digital rape” and an unethical use of Grindr. Verhoeven’s actions, he argued, violated the necessary “safe space” offered by gay-targeted social networking services like Grindr (Tsjeng, 2014).

Later that day, Verhoeven responded on his own Facebook page, framing Tilghman’s response as rooted in an unrealistic expectation of anonymity in networked contexts. Verhoeven argued that his project did not cross any ethical lines in publicly displaying Tilghman’s profile and conversations with the artist because, in Verhoeven’s words, anonymity on Grindr is “a myth.”

¹⁴ <http://www.driesverhoeven.com/en/project/wanna-play>

Everyone who loads Grindr or a comparative app. on their smartphone can see the photos and profiles. In the agreement with Grindr users have to accept that their information will also be viewable without having to be registered (from the agreement: “You acknowledge that some of the Grindr Services may be accessed...without the need to register an account.”).¹⁵

Verhoeven positioned his project as a use of publicly-available information on Grindr, which he then manipulated and redisplayed. Nothing in his installation, he suggested, did anything that any interested party with a smartphone couldn't accomplish himself. As a solution to any potential concerns about privacy, he invited men who didn't want to participate in the project to block him on Grindr.

Grindr's staff disagreed with Verhoeven's assessment of the ethics of the installation, characterizing his interactions with other users as “entrapment” and urging users to “flag” Verhoeven's profiles to bring them to the service's administrators' attention. Grindr promised to ban Verhoeven. Following the outcry, Verhoeven logged out of Grindr on all five phones and put up a curtain in the exhibition space, obscuring his actions. Two days later, only five days into the planned fifteen, Verhoeven shut down “Wanna Play?”, acknowledging pressure from Berlin's gay community. Later that week, *The Guardian* reflected on the installation, summing it up with, “It is that timeless philosophical question: if everyone sees your dick pic hanging in a gallery except you, is it art?” (Cain, 2014).

Following the project's end, Verhoeven published a reflection on “Wanna Play?” on his website in English, German, and Dutch (Verhoeven, 2014). In it, he apologized to Tilghman for the harm his actions may have caused, and acknowledged the limitations of

¹⁵ <https://www.facebook.com/driesverhoeven/posts/736549573049596>

his system for obfuscating the personal information of other Grindr users. He conceded that explicitly discussing the nature of his project in interactions with users would have been the “morally correct” approach. But he insisted that the outcry over privacy was a red herring:

Another feeling crept over me: that this was no longer a protest for protecting privacy, but that there was a deep-rooted desire to make the whole phenomenon [of gay sexuality] invisible. As if I was not allowed to inform the heteronormative outside world of the existence of the online cruising area, which is what Grindr is. I found it striking that 30 years after the call for the visibility of the homosexual community, there was now a protest for its invisibility. (Verhoeven, 2014, p. 2)

He closed the reflection with the hope that his project could prompt a frank discussion about the “substantive implications” of his actions — which, in his estimation, centered



Fig. 8: Public view of Dries Verhoeven’s “Wanna Play?” installation in Berlin, Germany. Photo by Sascha Weidner, courtesy of the artist’s website: <http://www.driesverhoeven.com/en/project/wanna-play>

on the persistent need to keep gay sexuality away from public view, even as public attitudes toward homosexuality have become increasingly liberal.

There is, undoubtedly, some truth to Verhoeven's interpretation of the outcry over "Wanna Play?" The management of public gay sexuality and the ethics of outing as a practice of forcible gay visibility remain fertile topics for academic and public engagement (Colter et al., 1996; Delany, 2001; Gross, 1993). But I want to insist on taking "Wanna Play?"'s critics (and the statements of the Grindr staff) seriously when they object to Verhoeven's use of the service as an invasion of individual privacy. The questions raised by users in this case are significant: What does the public resharing of profiles and conversations on Grindr mean for the men portrayed in them? What if, reasonably, some of the men featured in the installation are embarrassed by or angry about having their sexual interests or activities chronicled publicly? How can we conceptualize the use, reuse, and, potentially, misuse of individuals' personal information that's taking place? More generally, what happens when unexpected or unwanted things are done with people's data?

The problem with these questions, of course, is their vastness. Too often, they get reduced to black-and-white normative questions about the justifiability of a particular disclosure. The consensus interpretation of "Wanna Play?" is that Verhoeven behaved unethically because he didn't tell the men he spoke with that they were participants in a public art installation. But this normative criterion isn't generalizable beyond the specific issue of artists using Grindr in public, and it doesn't actually reveal anything about the underlying structures of data use and social acceptability that inform ethical judgements about this particular case. We can't, for example, generalize from the handling of this

case to the related practice of outing politicians who use Grindr (Bullock, 2011; M. E. Miller, 2015), or even to academics who publish accounts of their interactions on Grindr in peer-reviewed journals (Blackwell et al., 2014; Crooks, 2013; Raj, 2011). The specific cases we could examine are virtually limitless. The scale of the networked personal information ecosystem makes broad questions about the “rightness” or “wrongness” of sharing practically impossible to answer.

I want to propose a narrower, but potentially more instructive, set of questions: What do “unexpected” and “unwanted” actually mean? From whose perspectives should we be asking and answering these questions? Individual users? Men who have sex with men? The public at large? Business owners? Members of vulnerable populations that we believe need to be legally protected? Each clearly has something at stake in the networked afterlife of personal information; but it’s often unclear how to parse the various overlapping interests and stakeholders that make up online personal information ecologies. More generally, how do the changing technologies of networked sociability augment these processes of data use and misuse? What are the affordances and constraints created by different technical systems for the circulation of personal information? And are these data-sharing practices a new invention of the digital age, or do they speak to pre-networked practices of sociability that simply manifest themselves in slightly different ways online?

This chapter begins to parse some of these issues, offering a broad framework for the circulation of personal information in networked contexts that can better describe these practices and their associated risks. I begin by outlining a working definition for networked risk that allows us to examine how sociotechnical processes can endanger

individuals and their data. I argue that, in networked systems, risk stems from a lack of individual control over one's own personal information. Using this definition, I establish a model for evaluating information sharing that distinguishes between what I call *inter-context* and *inter-platform data flows*. While many instances of data sharing take place across networked platforms, I suggest that not all such inter-platform flows necessarily imply a change in the social context of that information. Conversely, some uses of personal information within a single technical platform move data between several social contexts at once. Some flows of personal information may cross technical platforms and social contexts simultaneously. A context- and platform-sensitive analytic approach maps out the architectural properties of data flows, identifying which patterns of personal information use create tend to create which specific risks to individual safety online.

This approach prioritizes the structure of data flows over their content. Adjudicating the ethics, appropriateness, or actual risks of a particular instance of sharing requires us to examine the substance of the data flow: what's being shared, about whom, and with what likely or actual real-world consequences. The infinite variability of networked information makes it difficult to translate these particularities into a generalizable model of data sharing. Instead, this approach examines structural characteristics that are common to many flows of personal information. This reveals risky classes of data flows, whether or not instantiations of those classes have already proven themselves to be damaging to particular individuals in particular cases. The goal of this meta-analytics of data flows is to enable users, software developers, and policymakers to identify appropriate, effective, and narrowly-tailored solutions that minimize risk in networked interactions. This allows us to move beyond ad-hoc responses to individual

instances of problematic data use and toward generalizable answers that can protect individuals in a wide range of circumstances.

The approach outlined in this chapter also moves away from the traditional analytic vocabulary of networked “privacy,” in favor of an approach focusing on individual data sovereignty. The dictionary definition of privacy centers around a freedom from observation — by other individuals, by governments, or, perhaps, by sophisticated electronic systems. Herein, privacy becomes, simply, an absence of visibility: a protection from observation enabled by keeping sensitive information out of public view. While helpful, privacy remains a limited analytic tool, ill-suited to account for situations in which visibility is desirable, or individually desired. By contrast, data sovereignty prioritizes individual control over the privateness or publicness of their information. A sovereignty approach takes as its end goal the establishment of sufficient systems to enable interested individuals to make appropriate choices for themselves about the risks they want to take when sharing their personal information. Privacy (or invisibility) is but one choice among many; individuals should have the ability to make these decisions in an active, informed manner.

Using this framework, we can focus on cultivating social and technical structures that minimize the potential for negative safety outcomes and maximize the possibilities for data sovereignty for people online. I illustrate this approach with two case studies of gay-targeted online services:

1. In-app advertising and the commercial structure of gay geosocial networking applications;

2. The blog Douchebags of Grindr, which publicly posts screenshots of Grindr profiles that its authors deem offensive or inappropriate.

In each case, personal information is used by an agent other than the data's original creator, for purposes that differ from the original goal of enabling networked social interactions between gay men. But, critically, the structural characteristics of these data flows differ in fundamental ways — as do their corresponding risks. The monetization of gay data moves information within a platform, but out of its original social context.

Douchebags of Grindr shares personal information across networked platforms while remaining (largely, though problematically) within a gay social context. Drawing on social, legal, and technical remedies, I argue that each class of data flow — intra-platform and intra-context — requires a differently-tailored, but ultimately generalizable, solution.

Data's risky afterlife

The previous two chapters have discussed the processes and practices by which gay men's personal information — what I've termed "gay data" — is gathered, shared, used, and managed. This discussion has focused on what we might term ordinary uses of data: applications of personal information that correspond with a reasonable user's reasonable expectation of what their data is going to be used for. For example, we could say that an ordinary use of a person's photo on Grindr is to share information about one's appearance for the purposes of participating in the in-app Grindr community and meeting other people. These uses are intuitive, obvious, and everyday.

By contrast, what I term data's *afterlife* focuses on off-label uses of personal information. This perspective takes the vibrant life of gay data, discussed in the previous

two chapters, as its point of departure; it brackets the contested and subtle forms of self-expression taking place on gay social networks as the reasonable, expected, quotidian interplay of individuals, communities, and technosocial systems. Instead, data's afterlife looks to an unknown and uncertain future for data, in which we find unanticipated (or perhaps unwelcome) actors and actions in the web of networked gay sociability. Talking about an afterlife for data recognizes that these potential unexpected, unwelcome, or badly-behaved actors are still intimately tied to the everyday actors and uses of information that we know, expect, and explicitly authorize. But it points to something about these uses of data that's less intuitive, obvious, and predictable. It recognizes that certain uses of data can make us anxious, or seem ethically wrong, or create risks of embarrassment or material harm.

Data's afterlife is the product of several basic properties of networked information. Borrowing from danah boyd, we can describe these characteristics as networked data's persistence, visibility, spreadability, and searchability (boyd, 2014a, p. 11). These affordances of networked information are well-documented, and we can trace many of the popularly-discussed anxieties about networked media back to some combination of them (Ambrose, 2013; Halavais, 2009; Mayer-Schönberger, 2009; Nissenbaum, 2010; Shein, 2013; Solove, 2006; 2008). Where a more nuanced definition is needed is around the concept of risk: that is, around those properties of networked media which create the potential for harms to individuals or communities.

I argue that, in networked systems, risk stems from a lack of individual control over one's own data. This ideal is encapsulated in the emerging framework of data sovereignty (Obar, 2013). This notion, implicit in market-based analyses of personal

information (e.g. P. M. Schwartz, 2004) and explicit in recent regulatory proposals like the Consumer Privacy Bill of Rights (White House, 2012), suggests that individuals should be empowered to control the creation, collection, and circulation of their personal information online in a granular and sophisticated manner. Most recently, the establishment of a formal “right to be forgotten” in Article 17 of the 2012 European Data Protection Regulation concretized these ideas, giving individuals the legally-sanctioned ability to secure the erasure of their personal information from websites operating within the European Union (Ambrose, 2013). Asserting one’s right to be forgotten is an act of data sovereignty. The logic of these proposals and actions is clear: The absence of adequate information and tools for informational control puts individuals at risk — of diminished personal sovereignty, and therefore of material harms to one’s reputation and professional life.

It’s worth pausing here to excavate two of the basic consequences of this formulation of online risk: First, that virtually all networked information is enshrouded in an aura of “being risky,” whatever its actual content or uses; and second, that framing the circulation of information as an individual risk positions data management as part of a neoliberal care of the self.

The canonic articulation of risk as an analytic object emerges in Ulrich Beck’s *Risk Society* (1992). Modern societies, Beck suggests, have transitioned from a focus on the social distribution of goods in an economy to the social distribution of risk among individuals and institutions. The distribution of material goods does not cease to exist; but it is no longer the central object of public concern. Risks, Beck argues, are the primary product of industrial society. Accordingly, like Anthony Giddens (1990), Beck suggests

that modern societies have become future-oriented, focusing on hazards and insecurities that do not yet materially exist but which we assume *will exist soon*. Embedded in these arguments is a sense of inevitability: that risks are salient to us because we know that their unpleasant futures are not just possible (in a dim, distant way), but are probable and therefore actionable.

This logic of risk applies broadly in the context of networked media. It is impossible to conclusively declare that any given piece of information will never be embarrassing or discrediting or undesirable in any situation we might encounter. The increasing persistence of data extends this unknown horizon indefinitely into the future. The basic properties of networked information render it as constantly at risk of being used in an unexpected, unauthorized, and damaging fashion, by individuals or in situations we neither know nor can reasonably predict. All data, in this sense, is inherently risky, because its future — what I call its afterlife — is both unknown and unknowable at the moment of its creation. When we choose to share data online, we do so in tacit acceptance of the fact that the benefits of sharing outweigh its unknown and as-yet-unrealized risks.

The problem herein is how to assign responsibility for data's unknown and potentially risk-laden afterlife. When we translate the omnipresent riskiness of data into a call for individual data sovereignty, we position data as a site of the Foucauldian care of the self (Foucault, 1986). By declaring data sovereignty a "right," governments establish data as a site of actionable risk and concern: that there's something about information that justifies a special empowerment of individuals. Having sovereignty over one's data means assuming the responsibility for that data. In classic neoliberal fashion,

governments are able to act at a distance (T. Miller, 1993), making individual autonomy the basis for the protection of personal reputation and welfare. This has two effects. First, it places responsibility for any failures of data management on the individual, rather than on institutions or governments. This absolves governments and service providers of the need to care for individual users; they need only provide the infrastructure for people to take care of themselves. Second, it creates yet another arguably uncompensated class of labor for individuals whose data circulates online (Andrejevic, 2011; Terranova, 2000). Providing individuals with the tools to take action places the burden to *actually* take that action squarely upon their shoulders. The management of as-yet-unknown risks becomes a task that requires constant vigilance from individuals who want to assert control over their networked identities.

Despite these concerns, I argue that data sovereignty represents the best available framework for balancing risk against the needs of other people and communities in data's networked afterlife. It's impossible — and, I would argue, undesirable — to completely eliminate risk in networked interactions. To do so would close off possibilities for new or novel forms of interpersonal interaction across emerging media — innovative practices which always imply a degree of uncertainty and potential harm. Nevertheless, I agree with the broad prescriptions of a data sovereignty approach that argues that individuals have a basic right to control over their personal information, inasmuch as that information can be understood as a networked extension of the self. Empowering individuals to manage risk gives them a greater capacity to, in Daniel Solove's words, "be who they want to be" (Solove, 2008, p. 4).

Striking the correct balance between sovereignty and unfettered innovative interactions requires us to develop better solutions to the potential harms of risky data flows. I suggest that there are concrete steps we can and should take to make the contours of networked risk more clearly visible to people, and to make it harder for those risks to actualize into harms. Nevertheless, totalizing solutions that privilege unfettered data sovereignty do not reflect the actual subtleties of data's afterlife. In order to craft better solutions, we need to develop a granular approach to tracing the actual risks in the networked flows of data's afterlife.

Platform and context specific data flows: A model

Social information has a curious habit of refusing to stay where individuals put it. Tiziana Terranova has described this as “a tendency of informational flows to spill over from whatever network they are circulating in and hence to escape the narrowness of the channel and to open up to a larger milieu” (Terranova, 2004, p. 2). Terranova's description outlines the stakes for this chapter's approach: How can we precisely describe the “tendencies” of networked information flows that define data's afterlife? I offer a granular analytics of technical platforms and social contexts as a solution. My goal herein, following Terranova's call for inquiry into what she terms “informational dynamics,” is to provide a suitable language for describing the networks, channels, milieus, spillages, and escapes that we experience as data's afterlife. This section maps out the basic terms of my approach.

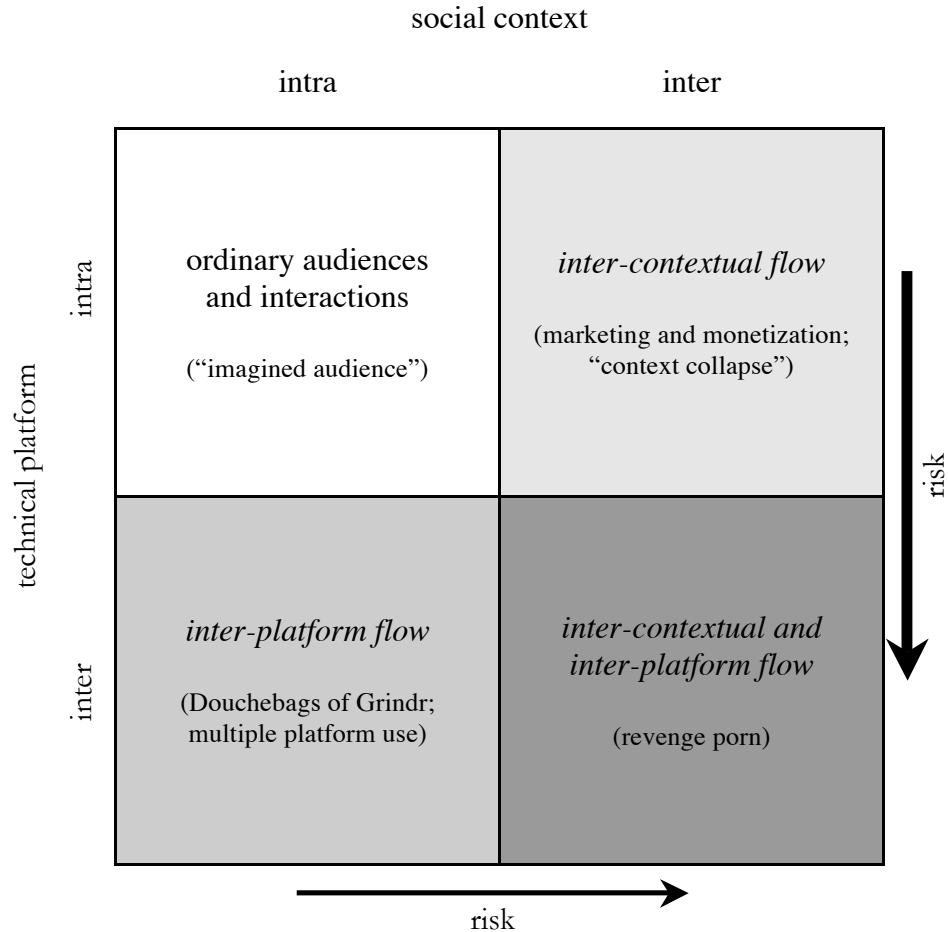


Fig. 9: Data flow model. The greater the change in platform or context, the higher the risk. Flows across both context and platform pose the highest risk to user safety. Flows across platform are riskier than flows across context, because they may (intentionally or unexpectedly) also become inter-contextual flows.

I identify two basic units of analysis that we can apply to any networked data flow: *how data is managed* (technically) and *what data is for* (socially). We can define two variables to represent these two properties: the *platform* that contains it, and the *context* in which it is used. The default conditions for these variables correspond with the initial state of data immediately following its creation by a technical system or disclosure by an individual. These conditions can include any number of specific attributes that

describe a platform or context: who has access to information; what format it's in; how it can be exported or shared from a given system; what the rules of proper interpersonal conduct are surrounding the use of that data; and so on. Any piece of information can have its status within a sociotechnical system described along these two axes.

Ordinary data use keeps information within its initial platform and context. For example, a user who creates a profile on Grindr expects that information to be used (1) within the Grindr application, (2) by other Grindr users, and (3) with the purpose of facilitating social connections among Grindr users. The details of these ordinary uses may vary from person to person, but the general properties and affordances of a service or piece of software suggest certain standard uses. Grindr is a gay-targeted social networking service; it's manifest function is to enable networked interactions between gay men. We can define this as the initial platform and context for information shared on Grindr.

When data flows through networked systems, we can describe that flow as potentially affecting either data's platform or its context — or both. Not all such flows share the same properties. Information may cross contexts without changing platforms; or it may change platforms without dramatically changing social context. Neither of these flows necessarily imply a corresponding change in the content of information; context and platform describe the paratextual wrappers that surround information. They describe where, how, and under what circumstances information is used. A change in either platform or context may cause a change in the perceived or actual risks associated with a piece of information, even if the content in question has not changed.

Our goal in disaggregating platforms and contexts is to enable an analysis of data flows that recognizes the actual sources of risk in networked interactions, rather than merely the presence of upsetting or unwanted information exchanges. This granular perspective allows us to identify solutions to risk that directly address its underlying causes. I suggest that changes in platform may be most effectively counteracted with technical solutions that create systematized barriers to risky data flows. Risky changes in social context, on the other hand, can best be addressed by making users aware of how data is used, and implementing carefully-targeted policy solutions to legally proscribe especially risky inter-contextual flows. These solutions can be applied to a wide variety of cases — provided we can effectively identify what type of data flow is taking place in a given instance.

Commercial afterlives

First, I want to use a context-sensitive approach to data flows to map out the contours of a process familiar to virtually every user of a social networking service: the translation of personal information into commercially valuable consumer profiles. Through the use of cookies, persistent device identifiers, tracking pixels, and a host of other technologies, online service providers like Facebook and Google (as well as advertising networks and data brokers) are able to gather enormous amounts of information about individuals on the internet. In parallel, service providers are also able to mine information from the data that users themselves share — for example, turning a dating profile into a lucrative set of demographic data points.

Through an examination of how information flows across different social contexts within a single technical platform, I want to establish an analytics of risk for the commercialization of personal information on social networks. These practices of commercial profiling are commonplace in the domain of social networking services, but their ubiquity should not prevent us from critically examining their structure and consequences. I argue that the uses of personal information that enable commercial profiling constitute an inter-contextual flow of users' data. This section establishes an operational definition of inter-contextual data flows, then uses Grindr as a case study to demonstrate how this approach can be applied to the policies and practices of a particular social networking service. Ultimately, I suggest that, like all inter-contextual flows, the risks involved in the commercial use of personal information by a service provider can be mitigated through clear, accessible data use policies that enable users to make active choices about how, when, and where their personal information is used.

Crossing contexts

Often, people speak colloquially about things being “taken out of context.” This phrase has intuitive appeal in discussions of networked risk and privacy: It suggests that we share information within bounded social spaces with common rules of acceptable conduct, and that we face risks whenever information leaves those spaces in unauthorized ways. If information is decontextualized, it may be exposed to new, unexpected audiences, with unknown and potentially dangerous outcomes. Privacy — and by extension, individual autonomy and safety — is upheld by keeping things in context.

Helen Nissenbaum's *Privacy in Context* (2010) uses this idea of context as the premise for a broad framework for conceptualizing networked privacy. But what does "context," in an analytic sense, actually mean? At its core, the idea of "context" is a recognition that, over the course of our lives, we go through the world as actants within situationally variable and interrelated social systems, all of which have their own histories, practices, and rules. Nissenbaum offers this concise definition: "Contexts are structured social settings characterized by canonical activities, roles, relationships, power structures, norms (rules), and internal values (goals, ends, purposes)" (Nissenbaum, 2010, p. 132). Her definition has the benefit of accommodating virtually any attribute of any semi-structured social system as a potential determinant of context; but this permissiveness weakens the analytic utility of context as a tool for understanding information flows.

I agree with the conceptual core of Nissenbaum's overall argument in *Privacy in Context*: Different situations call for different (explicit and implicit) rules about information exchange, and we feel anxious about communication (online and off) when we believe that those rules haven't been respected. But, instead of taking "context" as a one-size-fits-all label for social principles (norms) and technical rules (transmission principles), I want to offer a more narrowly-tailored definition of "context" that focuses solely on the interpersonal information-sharing dynamics of individual and group interactions, setting aside for the moment the technical structures of networked data flows.

In my use of context, I'm drawing on performance theory, taken primarily from Erving Goffman's work in *The Presentation of Self in Everyday Life* (Goffman, 1959). In

chapter 2, I argued that the basic idea of performance offered by Goffman is insufficient to robustly describe how identity works in networked contexts. Instead, I outlined how surveillance, control, and confession represent a three-part model for understanding how individuals manage their identities in online spaces — both through active, conscious authorship, and automatic (technological or social) disclosure.

In this chapter, I want to engage with a yet narrower slice of Goffman’s performance theory: namely, the notion of a bounded social context within which performances take place. The principal characteristic of a social context that I draw from Goffman’s approach is the idea that individuals act within already-existing but infinitely variable socially-derived spheres of acceptable and unacceptable conduct. Put another way, there are discernible sets of rules that structure how social activities take place — and those rules have an impact on the choices individuals make about how to present themselves and interact with others. Individuals manage the circulation of their personal information (that is, others’ impressions of them) by carefully crafting performances of self that correspond with the norms of acceptable conduct within a given context.

Performance theory, as explained in *The Presentation of Self in Everyday Life*, establishes social context as the division between what Goffman calls “front stage” and “back stage” performances. The front stage describes the performance an individual intends to give: the traits that individual chooses to share, the manner in which she chooses to share them, her reasons for sharing those traits, and with whom she believes they’re being shared. Goffman describes this as a social space’s “working consensus”: the rules and social frameworks that describe a temporarily stable configuration of individuals and communicative norms that govern a particular interaction. The front stage

describes the performative context as individual performers understand it. In ideal performances, individuals are able to maintain a strict division between this front stage and the “back stage”: a category that describes performances of self that are still works in progress, or which aren’t intended for a given audience. A successful performance requires individuals to carefully control access to the back stage, and to effectively manage any disruptions to the division between the front and back stages.

In the event that a social space is disrupted — for instance, if back stage performances of self become available to the audience for a particular front stage performance — Goffman emphasizes that the individuals involved in a performance take active steps to resolve those disruptions and avoid their recurrence in the future. Humorous anecdotes and jokes, he suggests, make it clear what types of conduct are expected of individuals in a performance’s audience, and how any embarrassing disruptions of the performance should be resolved (Goffman, 1959, p. 7). What Goffman terms “civil inattention” also plays a significant part in this: becoming party to a performance of self that wasn’t intended for you should be resolved (under a functional social system) by acting as if you didn’t notice it (Goffman, 1959, p. 152). In this manner, both performers and audience members work collaboratively to keep performances properly situated within their intended social contexts.

Goffmanian performance describes an ideal type — a manner of sensitive, agential interaction between individuals and their interlocutors that can accommodate disruptions, provided everyone involved is willing to abide by the social norms of a given context. Goffman himself recognized the limitations of his work with reference to the boundedness of performative contexts. Writing in 1959 about Anglo-American societies,

Goffman notes, “We lead an indoor social life. We specialize in fixed settings, in keeping strangers out, and in giving the performer some privacy in which to prepare himself for the show” (Goffman, 1959, p. 157). But what happens, to follow the analogy, when our social lives take place outdoors? How can we account for the actual performances of self that take place in porous social contexts with sometimes ambiguous norms and rules?

A range of scholars have taken up where Goffman left off, interrogating the margins of performative contexts. In particular, the question of what Goffman calls “institutional integration” — the relationship among different, largely separate social establishments (Goffman, 1959, p. 153) — has proven to be an important area for research in studies of networked sociability. The term “context collapse” describes some of these interactions between otherwise separate social spheres (boyd, 2014a; Davis & Jurgenson, 2014; Duguay, 2014; Marwick & boyd, 2011; Wesch, 2010). Context collapse was initially described by Michael Wesch (2010) as a way to conceptualize the social operation of webcams in video performances of self. Individuals who share videos of themselves on YouTube, Wesch writes, do not do so into a networked ether, devoid of specific audiences and social cues; rather, webcam performances exist at the nexus of a potentially limitless number of different audiences and contexts — a state Wesch termed “context collapse”:

[Context collapse is] an infinite number of contexts collapsing upon one another in that single moment of recording. The images, actions, and words captured by the lens at any moment can be transported to anywhere on the planet and preserved (the performer must assume) for all time. The little glass lens becomes the gateway to a black hole sucking all of time and space — virtually all possible contexts — in on itself. (Wesch, 2010, p. 23)

Context collapse creates a crisis of self-presentation. In place of the neatly bounded front and back stage performances Goffman described, the condition of context collapse requires performers to reckon with some of the essential properties of networked information — its persistence, visibility, spreadability, and searchability, to return to boyd’s description (2014a, p. 11) — in the moment of self-presentation. The multiplicity of functions, informational norms, and potential audiences present in networked settings complicates notions of fixed social context.

An especially significant element of these discussions of context collapse is the idea of the audience. In Goffman’s framework, performers take their cues about what information to share, as well as how to share it, from an understanding of who the audience for their performance is going to be. Networked contexts, unsurprisingly, make it difficult to state with certainty what the audience for a particular piece of information actually is or will be. Nevertheless, individuals still have what Alice Marwick and danah boyd (2011) term an “imagined audience” for their communications: a group of people that are (implicitly and explicitly) addressed in the content of networked performances of self. These imagined audiences are structured by the manifest characteristics of networked spaces — for instance, one’s list of friends on Facebook or followers on Twitter — but also can include other individuals or groups whose specific identity is unknown. Anyone sharing information online, Marwick and boyd suggest, has some idea of the audience they’re addressing, even if they don’t know the specific identities of every individual in that audience. This idea of audience structures what individuals choose to share as well as how they share it, even if, in practice, the actual audience for information differs from its imagined one. But the fact that these audiences are imagined

rather than known complicates attempts to describe contexts in a bounded, Goffmanian sense.

I want to draw two key insights from this discussion: First, and straightforwardly, individuals often attempt to manage how they present themselves in interactions with others. To the greatest degree possible within a given social context, individuals try to control how others perceive them. This requires all parties involved to subscribe to a set of shared contextual norms of behavior. Second, and more importantly, this management is not always successful — or even possible. Performances may not go as planned, or audience members may not abide by the rules of a given context. Unexpected individuals might become party to performances not intended for them. In these cases, the relatively fixed and bounded social systems Goffman described in *The Presentation of Self in Everyday Life* are rendered unstable (the condition of “context collapse”), and therefore become risky to the individuals who interact within them. This riskiness can be exacerbated by technical systems that increase possibilities for access to a given social context while simultaneously making it harder for individuals and groups to enforce behavioral norms (Meyrowitz, 1985).

When describing data flows in and across networked contexts, we need to be sensitive to these possibilities for risk and contextual failure. Descriptions of data flows should begin by describing the Goffmanian ideal type of a performative context, then proceed to examine how the actual circulation of data might complicate the imagined boundedness of that context. To do this, we should begin by asking these questions about the social context of a particular networked interaction:

- Who are the manifest participants (audience) in the social context?

- What are the manifest social functions of information within the context?
- Who has access to information within the social context, whether or not their participation in the context is manifest?
- What are the formally described rules of proper information sharing behavior in the social context (for example, as articulated in the End User License Agreement or Terms of Service) as they govern all explicitly authorized participants?

Based on these traits, we should be able to conclusively answer the following question:

Does a given data flow — either in terms of its participants, its rules, or its functions — augment the manifest properties of a social context?

If the manifest functions, explicit rules, or norms of conduct differ across the two endpoints of a data flow, we can define it as an inter-contextual flow. These flows reflect qualitatively different understandings on the parts of the actors involved about what networked information is intended to be used for. In these situations, risk emerges as a result of differences in data's manifest and actual uses. This manifest difference represents a change in data's context, whether or not the end use of that data is formally authorized by a service's operators. The degree of risk created by inter-contextual data flows varies depending on the degree of difference (in functions, rules, and norms) between the origin and destination contexts. Users may believe (correctly or incorrectly) that their information is “for” a particular purpose; but other actors with access to their information may put this data to a different and potentially unanticipated use.

While they are enacted through information systems, inter-contextual data flows are fundamentally social processes. Actors (be they individual users or commercial entities) choose to use data they have access to for a wide range of purposes.

Accordingly, I suggest that the most appropriate way to ameliorate the risks associated with inter-contextual flows is to make potential contextual differences as visible as possible to the individuals who use networked systems. This can be achieved through active user education and clearly articulated data use policies. Visible contextual boundaries enable individuals to make active, informed choices about the information they share (or elect not to share) in particular networked contexts, with an awareness of the likely participants in the social context (or contexts) within a given platform. When users are made aware of the polyvalence of their personal information, the risks posed by many inter-contextual flows (such as the commercial data flows I discuss in this case study) are relatively minor. Higher-risk inter-contextual flows should be addressed through both education and enforced policy guidelines about the permitted uses of personal information.

Paying to be sold

On Grindr, personal information moves between two discrete social contexts: data's manifest social context on Grindr (interpersonal interaction) and the context of commercial activity (targeted advertising). It's important to recognize that, while marketing and market-creation are built into the information structure of the Grindr service (see chapter 2), a majority of the service's revenue comes from paid subscriptions from individual users (Crook, 2013). This creates a point of tension between how Grindr says it treats its users' personal information and how it actually behaves.

The Grindr application is available in two versions: a free-to-use, advertisement-supported app, and a paid version titled "Grindr Xtra." The free version of Grindr

provides full access to all the basic elements of the service: creating and sharing a profile, viewing the profiles of other users, initiating and receiving chat messages, and exchanging photos or location data are all possible without users needing to invest any money upfront into the service. Yet, in the case of Grindr, the paid version of the app — which costs \$0.99 to download, in addition to a \$11.99 monthly subscription fee — has consistently made up the majority of the service’s revenue stream; recent figures suggest that as much as 75 percent of the service’s revenue comes from subscriptions, rather than advertising. These figures beg two questions: (1) Why offer a free version at all?; and (2) Why do users choose to pay for services they could access for free?

Particularly as the market for mobile applications has become more crowded with roughly equivalent services, would-be developers are forced to contend with the basic Web 2.0 dilemma of scale: a social service predicated on interpersonal interactions is only as useful as the size of its audience, and users are unwilling to pay for services that can’t connect them to the people they’re interested in meeting. Embedded in this is the basic diffusion effect of critical mass (Rogers, 2003): Enough early adopters need be attracted to a service in order to demonstrate its basic viability. Interactive social services — and, particularly, services like Grindr that create social networks on the basis of geographic proximity between users — need to demonstrate to prospective users that they can find and interact with other users with sufficient ease and frequency to warrant using the service at all. At this stage, making a service available to new users for free is often an effective way to reach critical mass. A monetization strategy (if it exists) is secondary to creating a viable core network. Second, once a service has reached critical mass, developers need to capitalize on the value of their network to continue to attract new

users (Wang, Chin, & Wang, 2011). Each new user past the point of critical mass contributes to the overall value of the social network, thereby reinforcing its position in the marketplace.

The need to achieve critical mass explains the importance of free versions of paid social services: even if a service's primary business model centers on paid customer relationships, it first needs to demonstrate to potential customers that a service is actually worth paying for. A *freemium* business model that makes the basic components of a service available for free, even as the full interactive experience is only made available to paying customers, enables service providers to cultivate and demonstrate the value of their network to otherwise skeptical potential users.

Among users who chose to pay for Grindr, their chief reasons for subscribing centered around improving the in-app experience. In particular, removing visible advertisements, increasing the number of potential connections, and accessing additional search and interaction features were among the most frequently cited explanations for paying an app's subscription fee. Many users cited convenience and ease of use as the motivating factors behind buying an app, emphasizing features that enabled them to automate routine actions like sending photos.

Significantly, none of the users who reported paying for an app claimed that the free version was patently unusable. Rather, paid versions are understood by users to offer premium improvements on the in-app experience that warrant, in some instances, an upfront investment. Users recognized the development choices behind the free/paid split as derived from business needs: As one user put it,

It's the small conveniences. ... Basically, [paying for an app gives you] the experience you'd expect from a good app. That they [the developers] purposely made [worse] in the free version. In paying, you expect to get the "best version" of the experience. But it's a fallacy sustained by the flaws and limits put in place in the free version. (Y. Roth, personal communication, January 25, 2014)

Despite complaints about the limitations of free versions, users did not indicate that paying for an app improved the quality of the in-app network itself or the core social experience of using an app. To the contrary, several respondents posed counterfactual scenarios in which only a paid version was available, subsequently noting that the result would be fewer potential interactions and therefore a diminished overall user experience. The paid version of the application is understood by users to be the urtext from which the free version is derived, even as the viability of the paid version depends on the existence of its free counterpart.

In addition to the in-app experience, users noted that paying for an application changed their expectations about the levels of support and reliability they receive from the developers themselves. Both in interviews and in reviews of Grindr available through the Apple App Store and Google Play, users complained about intermittent service outages and bugs in the client application. In articulating their dissatisfaction, users of the paid version emphasized their status as paying customers, noting that glitches and service outages "diminish the value of [a] subscription app." Several of the users interviewed indicated that they stopped subscribing to the paid version of an app on the basis of technical issues; in the words of one user, "If I pay for a service, I expect to be able to use it. When that doesn't happen, I stop paying. Nothing else matters when you can't use the app" (Y. Roth, personal communication, December 9, 2013).

Several users cited a sense of social responsibility to independent application developers as a motivating factor in their decision to pay for an app. As one user put it, “I like to support products I enjoy” (Y. Roth, personal communication, January 12, 2014). While a feeling of obligation to developers was not a dominant discourse among respondents, it highlights an important property of the mobile application economy: many of the developers creating mobile applications — including popular apps like Grindr and Scruff — are much smaller-scale businesses than social networks like Facebook and Twitter. Particularly in the case of gay-targeted services like Grindr and Scruff, where the hypothetical maximum size of the app’s user base is orders of magnitude smaller than that of mainstream social services, some users understand themselves to have a more direct and personal relationship with developers. This increases the expectations those users have for reliability and responsiveness to support requests; but it also establishes a stronger first-party relationship between developers and users. The prominence of particular figures associated with each app — for example, Grindr’s CEO Joel Simkhai and Scruff’s founder Johnny Skandros — reinforces the sense that users are investing in the humans behind services they value, rather than corporations interested solely in monetization. In contrast to the “black box” of social platforms like Facebook (Gillespie, 2010; 2012; Turow, 2012), services like Grindr and Scruff appear to be more directly connected to their users by virtue of the status of their founders within preexisting circles of gay sociability.

Even as nearly all the users interviewed identified the removal of in-app advertisements as a contributing factor in their decision to pay for the Grindr service, they tended to refer to advertisements as issues of aesthetics or convenience. In-app



Fig. 10: Illustration showing the placement of advertisements in the Grindr interface, as well as suggested use cases for Grindr marketing. Photo courtesy of Grindr, available at <http://grindr.com/advertise>

advertisements, one respondent noted, “take up as much space as a row of four pictures. They make using [the app] a worse experience” (Y. Roth, personal communication, January 20, 2014). None of the users interviewed identified targeted advertising (or the information from their profiles that is used to target advertising) as a concern, until they were specifically prompted to discuss the issue of privacy. This framing is significant: Considering advertisements as aesthetic or usability concerns calls into question one of the seeming truisms of social media: that users object to tailored advertisements online on the basis of privacy concerns. In the case of Grindr and Scruff, users objected to advertisements because they found them visually and interactively annoying.

Despite users’ overall attitudes of nonchalance, we should recognize a persistent discordance between Grindr’s rhetoric about user privacy and its actual treatment of personal information. On one hand, Grindr consistently stresses that its business model is

built around protecting the privacy of its users. For example, in response to publicly voiced concerns about the security of location data shared through Grindr (Mowlabocus, 2014a), a post on the Grindr blog stressed that user safety — defined here, principally, as the protection of location information from being accessed outside of the Grindr app by unauthorized parties — is a paramount concern for the service’s developers: “There is nothing that matters to us more than the safety and security of our user[s] and the Grindr community. We will continue to find ways to keep our users private” (Grindr, 2014c). Grindr’s CEO has likewise emphasized that, whenever possible, the service makes attempts to eschew collecting sensitive private data from its users: in his words, “We just don’t keep that kind of information” (Erlichmann, 2012). And, more generally, public statements from members of the Grindr staff emphasize that the service was built to enable interpersonal, rather than commercial, interactions: “We are always focused on doing what we’ve set out to do from the beginning: help guys meet other guys” (Grindr, 2014b).

These statements may be sincere, but they speak to a tension at the heart of Grindr’s status as both a social and commercial service. Protecting the privacy and safety of users is a sometimes contradictory goal to enabling a business model built in part around highly targeted advertising. The successful paid “Xtra” service suggests that much of Grindr’s revenue stream does not rely on the monetization of personal information; instead, Grindr makes a useful service available to its users for a monthly fee. This differs dramatically from the free-to-use financial logic of mainstream services like Twitter and Facebook. Despite this, and despite numerous statements from the service’s developers emphasizing the importance of user privacy, Grindr maintains the infrastructure to mine

its users' data for marketing purposes in both the free and paid versions of the app. Even users who pay for the Grindr service are not exempted from the market-making practices that enable lucrative targeted ads in the free version of the app.

Grindr's privacy policy makes it clear that the service reserves the right to employ users' personal information for a wide range of purposes, including marketing. The privacy policy is only visible once in a user's experience on Grindr, in a lengthy scrolling click-wrap privacy policy that users are forced to accept the first time they launch the app. The policy, while extensive and often detailed, gives the service's developers significant latitude in their uses of user data. For example, the policy gives Grindr the ability to use an individual's profile information, defined as

[a user's] photo, display name, status, relationship, looking for, ethnicity, age or date of birth, geo-location data, email address, password for the Grindr Services, height, weight, social network link, "Favorites", "Blocks", "Tribes" and any other information [a user] voluntarily add[s] to [his] profile on the Grindr App or is generated by [his] use of the Grindr Service (Grindr, 2013a)

for any of several purposes, including "to provide the services [users] request" — which, for the purposes of the privacy policy, includes serving advertisements within the Grindr application interface.¹⁶ In this sense, Grindr's use of personal information for marketing purposes is legitimate by the letter of the service's privacy policy. But the ambiguities created by these dual revenue streams raise significant concerns about whether Grindr does enough to make its treatment of users as sources of market data visible to all users of the app — including those who pay for the Grindr Xtra service.

¹⁶ In this context, users "request" the display of advertisements when the Grindr client software installed on their device attempts to pull the ad content from Grindr's servers or the servers of Grindr's ad networks. "Request" does not necessarily imply that users want or endorse the advertisements themselves.

Solutions: Contextual disclosure

The essential questions resulting from this discussion are: Is Grindr's use of personal information for marketing purposes deceptive, based on the tensions between its stated aims and its actual uses of user data? And, whether or not they are deceptive, do ambiguities about the circulation of user data on Grindr constitute a source of risk for individuals?

In order to answer these questions, we need to map out the technical and social information architectures of data flows within Grindr. In this case, data remains within the bounded technical limits of the Grindr platform; information does not circulate to other networked services, or outside of the control of either Grindr users or the Grindr staff. Information does, however, appear to move between two discrete social contexts: interpersonal interaction and marketing. In order to more specifically describe these contexts, I return to the framework for inter-contextual data flows outlined earlier:

- Who are the manifest participants (audience) in the social context?

The only participants overly revealed through the Grindr app are other users.

- What are the manifest social functions of information within the context?

Grindr markets itself as a social networking service built by gay men for gay men, emphasizing interpersonal interactions as the *raison d'être* for the application. In this instance, targeted advertising is an additional, though not always manifest, use of personal information.

- Who has access to information within the social context, whether or not their participation in the context is manifest?

Two groups of individuals have access to information on Grindr: (1) Grindr users themselves, who can access profile information for up to 250 individuals in their geographic proximity, as well as information about any individuals with whom they are actively communicating; and (2) the Grindr staff, including ad sales representatives.

- What are the formally described rules of proper information sharing behavior in the social context as they govern all explicitly authorized participants?

The final question highlights the potential sources of risk in commercial data flows on Grindr. Based on a reading of the service's privacy policy, we need to determine whether Grindr gives its users a reasonable expectation that the context within which their information is used also includes its use for marketing purposes. In practice, I argue that this is not clearly the case and therefore represents a source of risk to users.

Despite the detailed information on the Grindr ad sales website, the app's privacy policy does not explicitly list targeted marketing as a use of personal information. The information about targeting techniques that is readily available to interested marketers at <http://grindr.com/advertise> is not included in the privacy policy. Legally, service providers may not be required to spell out all the particular uses of personal information in a privacy policy; but general, categorical statements do not go far enough in establishing firm boundaries around the contexts for appropriate data use on Grindr.

Further, the rhetorical structure of the privacy policy often obscures the uses of information it is designed to describe. In an effort to define terms like "personal information" in significant detail, the policy sometimes offers contradictory descriptions of what information is actually used for. For example, the policy suggests that location information is only used for interactive purposes — "[A user's] last known location is

stored on our servers for the purpose of calculating Distance Information between [him] and other users” — even as the earlier definition of profile information includes “geo-location data” as a type of information that Grindr is permitted to gather and use for a broad range of purposes (including marketing). Based on Grindr’s advertiser guidelines, we can ascertain that location information is indeed used for marketing purposes; but the contradictory and confusing information in the privacy policy makes it difficult for even interested users to adequately determine how their location data is used.¹⁷

Inadequate information about the context for networked data constitutes a source of risk in the sense that it diminishes the ability of individuals to make informed choices about how, when, and with whom they want to share their personal information. A data sovereignty-maximizing approach to reducing networked risk would focus on clearly articulating the purposes for which personal information is used. In the case of Grindr, this would involve explicitly indicating that personal information shared in the context of using the Grindr service for its manifest purpose of interpersonal interaction is also, simultaneously, used to enable targeted marketing campaigns. A revised privacy or data use policy that makes explicit reference to marketing and advertising would be an important first step. More generally, however, I suggest that application developers should be forthright about the commercial structures of data use in the descriptions of their apps listed in the major mobile app distribution platforms (the Apple App Store and

¹⁷ It’s worth recognizing, of course, that contradictions, ambiguities, grammatical and typographic errors, dead links, and poor organization are common to many privacy policies. While I’m dwelling on particular problems with the text of Grindr’s policy documents in this dissertation, the fundamental issue is widespread, and undoubtedly contributes to the overall lack of comprehension of these policies demonstrated by Turow et al (2003). What’s needed is a philosophical change: from regarding privacy policies as legalistic formalities designed to comply with FTC requirements, to policy documents as clear articulations of the relationship between service providers and users that’s designed to empower users to make informed decisions about how their personal information is used and managed.

Google Play). In the case of Grindr, developers could indicate that all Grindr users, including those who pay for Grindr Xtra, are tracked within the Grindr app for marketing purposes.

These changes require very little additional labor on the part of app developers. Nor, ultimately, do they constrain the ability of developers to implement multiple revenue streams — as has quite lucratively been the case with Grindr. Nevertheless, more robust communication about the social functions of data has the important effect of giving users the information needed to establish a comprehensive model of how and by whom their information is used on Grindr — including in cases where the actual uses of data differ from the manifest functions of the Grindr app. Transparency on the part of service providers about the multiple contexts for personal information allows users to more actively participate in the management of their networked data, without unduly constraining the ability of service providers to explore a range of monetization strategies.

No fats, no femmes, no privacy?

Normatively speaking, many networked platforms are uncharted territory. For the early users of a new social networking service, it's often unclear what the boundaries of appropriate behavior within a particular context actually are. How are you expected to treat other users? What uses of a service's features are considered impolite? How should conflicts between users be resolved? These questions are often left unanswered by a service's developers, and are instead left to users to adjudicate for themselves. In a space of unknown or uncertain norms of conduct, it's often unclear how users should establish the ground rules for interacting with each other.

This portion of the chapter takes up the process of normative negotiation by examining gay men’s use of public blogs to capture, display, and discuss instances of perceived misbehavior on Grindr. Specifically, I look at the blog *Douchebags of Grindr*, which publishes screenshots of Grindr profiles that the site’s authors deem inappropriate, offensive, or otherwise “douchey.” Since its launch in 2011, the site has posted hundreds of profile screenshots, all of which display unobscured profile information — including, in many cases, a photograph that includes the user’s face. *Douchebags of Grindr* appears to serve two different functions: On one hand, the site gives Grindr users the opportunity to directly engage with the contested norms of proper behavior on gay social networking services — explicitly addressing instances of perceived racism, ageism, or “femmephobia” that otherwise only rarely enter into public discussion. On the other, the blog is a widely-read and widely linked-to source of entertainment, attracting a readership that is not limited to the men who use Grindr. In both cases, the public circulation of Grindr users’ personal information constitutes a source of risk.

Using a platform-sensitive approach, I argue that the risks of this off-label use of personal information result from an inter-platform flow of personal information: that is, from profile information moving between the relatively bounded Grindr platform and the widely accessible format of a public blog. I outline a careful approach to evaluating the structures and functions of data flows between Grindr and the *Douchebags of Grindr* blog: While the ongoing negotiation of behavioral norms taking place among Grindr users requires that “douchebags” be publicly identified as such, it’s difficult to deny that even accused douchebags ought to have an opportunity to know how and by whom their personal information is being used. Ultimately, I propose both policy- and software-

driven solutions that make it more challenging for users to remove other people's personal information from the Grindr platform without their knowledge. Adopting platform-enhancing solutions reduce the riskiness of invisible inter-platform flows, while not entirely foreclosing on possibilities for innovative multi-platform sociability. I suggest that any attempt to mitigate the potential harmful effects of inter-platform data flows need to recognize both their risks and, in many cases, their necessary social functions.

Crossing platforms

Recent scholarship about the internet has adopted the term “platform” as a way to describe a wide range of websites and applications. While the etymology of the term speaks to a wide range of offline practices — from raised surfaces to political agendas — “platform” has become a blanket term for the sociotechnical assemblage of code, policies, and users that defines an interactive networked service. A platform, notes Tarleton Gillespie (2010), is at once a way of referring to the computational infrastructure of a service and its sociocultural dynamics. This definition, now canonic, has become a central piece of many critical analyses of the social web (Bucher, 2012; Crawford & Lumby, 2013; Hands, 2013; van Dijck, 2013a; 2013b; van Dijck & Nieborg, 2009; van Dijck & Poell, 2013).

While helpful in many ways, Gillespie's insistence on examining the politics of platforms has obfuscated two critical dimensions of networked services: (1) Their technical properties; and (2) the points of integration between different platforms. The first concern is straightforward: In prioritizing the political properties of networked

platforms, we neglect their technical dimensions. Most platform-driven analyses focus on the political consequences of a platform's use, rather than on the granular structures of code and policy that give rise to those consequences. To counteract this tendency, we need a return to technicity in platform studies.

Second, “platform” is often used as a general category to contain individual networked services. Platform-driven analyses overwhelmingly tend to examine particular applications or websites in isolation, focusing on a given platform as an insular social and technical system. In my view, this does not adequately reflect the integrative realities of networked systems. Despite the rise of a relatively small number of major players, the internet remains a fundamentally fragmented space. While Facebook boasts the largest user base of any social networking service, competitor services like Twitter, LinkedIn, and Google+ nevertheless attract significant numbers of users by offering unique or specially-targeted features. Many researchers have used this as the impetus to chronicle the specific features or affordances for individual self-expression on different services, highlighting how factors like interface design influence what users are able to share, and with whom they are able to share it (Papacharissi, 2009; van Dijck, 2013b; 2013c). Users “have one identity,” notes José van Dijck (quoting Facebook's Mark Zuckerberg), but they constantly struggle to strategically mobilize that identity within the preexisting frameworks of available social networking services (van Dijck, 2013c). This analytic approach, which prioritizes the particularities and affordances of individual platforms, helpfully illuminates how individual self-expression varies between different networked settings. But it often leaves unexamined an equally important part of user behavior: namely, how individuals manage their identities *across* different social networking

services, using multiple platforms in parallel in different ways in order to accomplish their goals. We can no longer, as Nancy Baym (2011) has put it, conceptualize technical platforms in isolation. Rather, to paraphrase Henry Jenkins's discussion of transmedia storytelling practices in *Convergence Culture* (Jenkins, 2006, p. 96), we should imagine identities as horizontally integrated across multiple platforms, taking advantage of the specific affordances of each to contribute, in an aggregate sense, to an individual's overall networked presentation of self. What we need to interrogate, therefore, is how data flows through these horizontally integrated social systems. To do so, we can examine how information is managed on particular platforms, with an eye to how those information management practices either enable or prevent flows of data between platforms.

In order to develop better frameworks for describing the risks created by networked data flows, we need to disaggregate the technical processes of platform changes from the social processes of context breaches. Often, as in Nissenbaum's contextual integrity framework, the two merge and become a gestalt description of how data moves around sociotechnical systems. In this case, I want to propose a more technically specific vocabulary for understanding the circulation of personal information across platforms. In particular, I'm interested in the moment of translation whereby two platforms have to be made commensurable with each other, allowing information to travel between discrete websites or applications. Individuals take advantage of differing technical characteristics of networked platforms to enable them to share their personal information in order to achieve their expressive and interactive goals. But, importantly, these different technical characteristics often represent a barrier between services,

blocking the flow of information between platforms (whether users want that flow to take place or not). Whenever an inter-platform data flow takes place, an actor (or group of actors) is responsible for the act of translation between them. This act illuminates the potential sources of risk in an inter-platform data flow.

Accordingly, we need to ask questions that specifically identify the technical rules of data management in and between different platforms. One way of understanding these rules is to conceptualize them, following Alexander Galloway (2004), as protocols:

[Protocols] always operate at the level of coding — they encode packets of information so they may be transported; they code documents so they may be effectively parsed; they code communication so local devices may effectively communicate with foreign devices. Protocols are highly formal; that is, they encapsulate information inside a technically defined wrapper, while remaining relatively indifferent to the content of information contained within. (Galloway, 2004, pp. 7-8)

On the web, “protocol” typically refers to any of the commonly-used standards of networked information transfer: for example, the Transmission Control Protocol (TCP) and Internet Protocol (IP) networking models that, together, define the structure of the internet. But we can also, as Galloway does, interpret protocols more permissively, taking them as a general description of syntactic structures of information management.

Computer systems do not have an intuitive understanding of how to handle any arbitrary piece of information; they need protocols to explain the specific rules that govern data in the moment of its use. By studying these meta-structures of networked information, we can identify how particular networked platforms give rise to particular information-sharing practices — and, by extension, how on-the-ground uses of data might depart from the standard operating practices of a given platform.

We can readily identify, in a general sense, when information has changed platforms; it appears on a different website or service than the one on which it was originally created or shared. But, for the purposes of examining whether a given data flow involves a *risky* movement of information between platforms, we should begin by asking four questions about each endpoint of the data flow:

- What types of data are used on the platform?
- How is data gathered, stored, retrieved, and displayed on the platform?
- What are the steps required to export or remove data from the platform?
- What are the formal rules governing the storage and display of this information over the course of its lifecycle?

Based on these traits, we should be able to conclusively answer the following question:

Does a given data flow — either in the kind or quantity of data it handles, or the process by which it handles it — result in a change in the rules of data management between platforms?

Inter-platform flows take place when the structural characteristics of data storage and display are different at each endpoint of a data flow. The actor initiating a data flow has to reconcile the technical properties of the destination platform with the format of the data proved by its origin. If this act of translation modifies the architecture of information use and display, it constitutes an inter-platform data flow. This data flow creates a source of risk for individuals who shared their information within a specific system with the expectation that it would be governed by the technical rules of that system.

Significantly, this assessment of data management practices moves away from analyses focused on the manifest functions of a networked platform and toward a

protocol-driven analysis of networked information structures. Functionalist approaches to platforms blur the line between social structures of use and the underlying technical affordances and constraints of a given system. Instead of relying on categorical declarations of platform type — for example, “blog” or “social networking service” or “photo sharing application” — a technical approach to platforms identifies how information is, in practice, managed within a system and across different systems. This analysis allows us to readily identify whether or not those practices correspond with users’ reasonable expectations about how their data will be stored and displayed. This approach also reflects integrations between different platforms, rather than insisting on examining each platform as an insular sociotechnical space.

I want to emphasize both the utility and the risks of inter-platform data flows. On one hand, the ability to move information between networked platforms gives users the opportunity to actively and innovatively take advantage of the particular affordances of individual platforms to achieve their goals. Nevertheless, inter-platform data flows represent a significant source of risk in networked interactions — to which users are highly sensitive. The popularity of privacy-promoting applications like Snapchat, where user communications self-destruct after a specified period of time, suggests that many users anticipate and actively work to counteract inter-platform flows of their personal information (boyd, 2014b; Gillette, 2013). Of course, even privacy-promoting apps are fallible: taking a screenshot of an expiring image or showing your phone’s screen to someone else while a time-sensitive image is visible counteracts the technical structures implemented in Snapchat to protect user data for unauthorized sharing. This very fallibility highlights, for users and observers alike, that virtually all networked

information is subject to unexpected or unauthorized flows beyond the boundaries of a particular technical platform.

Inter-platform data flows are particularly risky because they tend to also become inter-contextual flows. When inter-platform flows take place — even when the participants in a context initiate those flows themselves — they frequently result in the increased availability of information that was once platform-specific. This increased availability may expose data to unanticipated or unwelcome audiences, causing, in essence, an inter-contextual flow (with all the consequences discussed above). This tendency for inter-platform flows to also become inter-contextual flows heightens the intrinsic riskiness of inter-platform data flows.

Ameliorating the risks posed by inter-platform data flows requires stronger articulations — in code, interface designs, and policies — of the boundaries between networked platforms. Strengthening the boundaries between platforms has the dual effect of making it harder for data to move unexpectedly between different networked systems and increasing social disincentives to initiate inter-platform flows in the first place.

“It’s just a preference”

There’s little room to dispute that many of the user profiles posted on Douchebags of Grindr are offensive. The several hundred posts on the site since its creation in 2011 showcase a wide range of creatively-expressed prejudices: racism, ageism, “femmephobia” (a dislike of perceived effeminate behavior), and “body fascism” are frequent tropes. A representative sample of posts on the blog includes profiles with captions that read:

- “Squinty eye, no reply.”
- “No bears, twinks, fats, fems. Please dnt ever hit me up if you will make me throw up with your pics!”
- “No Disease thank god! Im about to show You F A G s what youve been missing!”
- “I block ugly/old ppl that try to get at me.”
- “Dont even try talking to me”
- “Don’t be gay. Gentleman and Argentinians to the front. NO shorty’s, asians, fats, or fems. Be masc and funny.”
- “Hey what’s up looking for friends and people to talk too. Boxer briefs are a turn on. Sorry not into black people.”
- “no sushi aka no asian”
- “I’m a gay GUY! If I wanted to date someone feminine I would be straight and with a girl.”
- “I block more Asian than the Great Wall of china.”
- “hate everything. no fats, fems, olds, uglies or ethnics.”
- “Not in2 taste of Asia or India . If I wanted that I’d go to a restaurant.”
- “WHITES ONLY!! All blacks, keep moving cuz I ain’t interested unless u can prove not all blacks are the exact same mkay?”
- “What’s with Asians wanting to spoon? Don’t they use chopstix!? Ps: I’m not racist, I own a colour tv.”

- “Just a normal guy here. Not attracted to HIV+ guys, Muslims or Jews. No stalkers or anyone with a mental disability, either. I’m serious - NORMAL ONLY.”¹⁸

The profile screenshots are frequently accompanied by a caption explaining why the depicted user is a douchebag (for example, “OMG really racist douche”), as well as a survey that allows visitors to the site to vote (using a 1-5 scale, from “not such a douche” to “supermegadouche”) on precisely how objectionable they find the profile.

In part, *Douchebags of Grindr* speaks to distinctive historical negotiations of desirability and proper conduct that have substantial roots in American gay culture. Many of the themes chronicled on *Douchebags of Grindr* are instantiation of long-standing contestations of desirability in Western (and, particularly, American) gay culture (Han, 2007). On one hand, attraction and preference for particular sexual partners is a deeply individual process, rooted in personal histories, and highly specific combinations of social, cultural, and sexual contexts; one man’s Adonis might not get a second look from another. In the words of one commenter on the site, “And preferences are a bad thing because...?”¹⁹ Or, in a lengthier post from another visitor to the site,

REALLY?? so if someone is not into Asians or Blacks they are Racist?...it’s CALLED Freedom to Choose who you want and Don’t want Not Racist!!..Maybe if Blacks and Asians didnt pester the shit out of people on Grindr and other websites and could SPEAK proper English and Not Ebonics..They might get the respect they so Crave!!²⁰

¹⁸ I’ve chosen not to include screenshots from the *Douchebags of Grindr* blog that include Grindr profiles. Even if I were to obscure some personally-identifying details (such as faces) in the images, I’m not convinced that this approach would be sufficient to justify further publicizing the figures of the individuals in question. (The rather halfhearted tactic of blocking out only the eyes of a photo has been employed by other authors who reference publicly-available but reputationally-discrediting material. I reject the claim that this is sufficient, or indeed at all effective, in protecting the individuals portrayed.) While the information I reference is publicly available, I don’t want to participate in its further circulation. Interested readers can view the images I describe herein by visiting the *Douchebags of Grindr* blog.

¹⁹ User comment, posted on 9 June 2014 at <http://www.douchebagsofgrindr.com/2015/04/douche/>

²⁰ User comment, posted on 22 July 2011 at <http://www.douchebagsofgrindr.com/contact/>

Undoubtedly, people should have a “freedom to choose” their sexual partners. But, crucially, the poster’s reaction frames sweeping generalizations about groups of people — for instance, speaking in Ebonics — as reasonable, non-prejudicial expressions of preference. Further, any contact with users outside of the preferred groups is characterized as bothersome pestering. These justifications are common, and indicate the continuing contestation in gay male culture over norms of proper conduct and communication around sexual preference.

In practice, the expression of sexual preferences has tended to focus on predominantly negative sentiments about a relatively small set of markers of identity. Writing about text-based chats on Gay.com, Andil Gosine notes that disclosing his ethnic background (Indian) frequently resulted in abrupt conclusions to conversations that, moments earlier, had been engaged, lengthy, and flirtatious (Gosine, 2007, p. 144). Senthoran Raj (2011) has chronicled similar experiences on Grindr. And, across a range of networked services and national contexts, researchers have noted that Asian men are particularly frequent subjects of these discourses of preference (Riggs, 2012). This calls into question overly relativistic attempts to frame the logic of “no fats, no fems, no Asians” as “just a preference.”

The emergence of networked sites for gay interactivity has increased the visibility of this contestation of preference. In a study of gay men in Los Angeles, for example, Jay Paul, George Ayala, and Kyung-Hee Choi suggest that gay social networking sites have crystallized racialized sexual preferences, even though those preferences undoubtedly existed offline:

In face-to-face social interactions, race or ethnicity was a factor whose power was more often expressed in an oblique and coded manner and felt inferentially; online ads made clear and amplified the sense of race or ethnicity as a source of difference and value. (Paul, Ayala, & Choi, 2010)

They likewise suggest that the relative anonymity of social networking services is a critical affordance for this more overt disclosure of racial and ethnic preference. It isn't that the individuals expressing prejudicial or offensively-worded preferences are blithely unaware of the potential hurtfulness of their statements; rather, the authors suggest that individuals take advantage of the protections afforded by relative anonymity to express themselves more overtly (and therefore, potentially more offensively) than they otherwise might offline. In this sense, Douchebags of Grindr works to counteract the perceived pernicious effects of relative anonymity in networked spaces by preserving and publicizing the negative conduct of certain users.

More generally, it's worth bearing in mind that not all posts on Douchebags of Grindr are embroiled in this long-standing dispute over preference and desirability. While accusations of racism, ageism, and body fascism are among the most frequently-occurring themes in posts on Douchebags on Grindr, some posts on the site shame users for less obvious reasons. In November 2014, for example, the site's moderators posted a screenshot of a user's profile which included an unobscured face photo of a man in his early 20s with the caption "i am my hair." The screenshot also listed the city in which the user (presumably) resides. In the post's title, the user was branded "Hair Douche," and earned a rating of 3.43 out of 5 (with 34 percent of 161 anonymous voters branding him a "supermegadouche"). Another user's profile — which included both an unobscured face photo and a link to the user's Facebook account — was labeled "New Age Douche" for

including the caption, “Its [sic] all about the sound and the frequency at which you vibrate.” “New Age Douche” garnered a rating of 1.58, with a majority of voters indicating that the screenshot did not qualify the depicted user as a douche.

The tag cloud that appears in the right-side column of the website paints a more general picture of the blog’s content (see figure 11). For example, it includes a number of ambiguous designations for posts, such as “arrogant,” “hypermaterialist,” “hot mess,” “ugly,” “vapid,” and “weird.” These tags, assigned to posts by the site’s moderators, hint at the rhetorical instability (and therefore potential risks) of “douchebag” as a behavioral category. Unlike racism, ageism, body fascism, or even femmephobia — tendencies whose pernicious status within

gay communities has been widely acknowledged in both academic and popular literature — it’s unclear, on the basis of any descriptions posted to the site, why being merely “weird” warrants the same degree of public shaming as overt prejudice.

The ambiguities of the term “douchebag” point to the dual identities of the blog. On

TAGS

agism arrogant asshole assholes
at least he's not racist blockophiles body
nazi bottom cocky crazy cunts delusional dl
douche douchebag druggies dumb elitist
femmephobia grindr gross hair hop-
pers hippocrites homophobia hot mess hypermateri-
alism hypocrites idiot lolz masc megadouche
mess messy mezzy moron nipples racism
racist self-loathing supermegadouche tumblr
censors ugly unmedicated vapid weird

Fig. 11: List of recent tags used for posts on Douchebags of Grindr. Retrieved on 1 December 2014 from <http://www.douchebagsofgrindr.com>.

one hand, Douchebags of Grindr is a site of the active contestation of behavioral norms on gay social networking services. Men use the site as a venue for debating whether writing “no Asians” in one’s profile is a reasonable expression of sexual preference or an offensive comment that doesn’t belong on Grindr. Beyond the broad proscriptions offered in its Terms of Service, Grindr does not offer users any guidelines about appropriate conduct on the app; accordingly, users are left to themselves to establish new normative frameworks for the service, both within the app itself and across different platforms (including blogs like Douchebags of Grindr). As one visitor to Douchebags of Grindr put it, in an article about the site published on the Huffington Post: “The novelty of Grindr’s meeting space allows for a community-developed praxis. What speech is allowable as expression of sexual preference in this new queer space?” (Cooper, 2012). The answer, he suggests, remains open for discussion.

Critically, however, Douchebags of Grindr has also attained a degree of mainstream popularity within the genre of name-and-shame internet entertainment. The site has received frequent mentions on gay-targeted blogs and websites like Queerty and Towleroad, which describe Douchebags of Grindr as both a chronicle of all-too-familiar bad behavior on gay social networking sites and a source of comedy (Villarreal, 2011). More mainstream outlets have also taken notice, with mentions of Douchebags of Grindr appearing on the Huffington Post (Cooper, 2012; Whitney, 2012), Gawker (Moylan, 2011), and in the New Zealand Herald (Suckling, 2014). These mainstream mentions not only introduce new audiences to Douchebags of Grindr (beyond the initial circulation of the site in gay circles), but also reframes it as primarily comedic, rather than adjudicatory, in nature.

The publicness of the site also results in many of its posts appearing in search results for seemingly unrelated queries. A list of recent search terms which pointed to the blog, published in a sidebar on the home page, include: “tumblr naked men camping,” “naked kiwi men,” “naked icelandic men tumblr sites,” “naked arab men,” “naked korean men,” and “gay boy selfies.”²¹ The proprietary logic of the Google PageRank algorithm resulted in these queries returning links to posts on Douchebags of Grindr, despite the fact that none of them are related to Grindr at all. The site’s audience, therefore, extends far beyond only those Grindr users who are interested in adjudicating the platform’s behavioral norms.

Contested ethics

The dual nature of Douchebags of Grindr — of social tool on one hand, and public venue for entertainment on the other — has given rise to a broad contestation of the blog’s ethics. Even when readers do not find fault in the actions of the site’s authors, they recognize the potential harms inherent in the blog’s treatment of personal information; on balance, they simply find that the benefits of the site outweigh its risks. Other readers explicitly critique the techniques employed on Douchebags of Grindr, suggesting that the site may do more harm than good. I want to map out some of the debates in order to highlight the need for a less ethically contingent approach to flows of personal information.

I want to begin by stressing that, independent of its coverage in mainstream media outlets, Douchebags of Grindr is understood by its readers (and, especially, by Grindr users) as a venue for publicly adjudicating user misconduct that would otherwise go

²¹ Recent search terms included on the Douchebags of Grindr homepage on 16 April 2015.

unaddressed in the app itself. In part, Douchebags of Grindr exists because of a failure by Grindr's staff to appropriately enforce the service's articulated Terms of Service.

Alongside a range of standard prohibitions on the use of the Grindr service for unlawful purposes or the dissemination of unsolicited spam messages, the Grindr Terms of Service forbid users to

post, store, send, transmit, or disseminate any information or material which a reasonable person could deem to be *objectionable, defamatory, libelous, offensive, obscene, indecent, pornographic, harassing, threatening, embarrassing, distressing, vulgar, hateful, racially or ethnically or otherwise offensive to any group or individual*, intentionally misleading, false, or otherwise inappropriate, regardless of whether this material or its dissemination is unlawful. (Grindr, 2014a emphasis added)

A reasonable interpretation of this clause in the Terms of Service would suggest that many of the prejudicial statements chronicled on Douchebags of Grindr shouldn't have been permitted to appear on the service in the first place. While Grindr users are able to report particular profiles as offensive within the app itself, the frequency with which these profiles appear suggest that either Grindr is unable to effectively enforce this portion of its TOS, or that the service's operators are unwilling to wade into the "just a preference" disputes.

In an important sense, therefore, Douchebags of Grindr is a user-derived intervention into a perceived failure on the part of a service provider to effectively and sufficiently govern user conduct. A number of appreciative comments posted on the site reflect the perceived social need the site fulfills:

Thanks so much for creating this website. The widespread douchebaggery on Grindr (not to mention manhunt, a4a [Adam4Adam], etc.) is disturbing, sad, and ugly. Shining a light on gay male racism, body fascism, and

assholism is definitely needed and this is a great start. Maybe these guys will actually reflect a bit and see the error of their ways.²²

This commenter positions Douchebags of Grindr as a powerful actant in normative negotiations around gay online conduct. The individuals portrayed on the site, the commenter suggests, need to be shown “the error of their ways.” This posits a causal link between online shaming and behavioral modification: being called out as a douchebag of Grindr is presumed to have the effect of causing individuals to change how they express themselves on the service. Douchebags of Grindr is thus framed as an important instrument of behavioral intervention, rather than merely as a source of comedy for its readers.

More frequently, endorsement of the site are qualified with objections to the tone, style, or approach the blog’s moderators adopt. In the words of one commenter, “So you’re pointing out how people are assholes by being bigger assholes yourselves?”²³ Comments of this type tend to focus on the titles, captions, and tags attached to the posts, rather than the profile screenshots themselves. The commenters do not dispute that many of the individuals portrayed on the site are violating behavioral norms; rather, they suggest that the shaming tactics employed on Douchebags of Grindr may not be dramatically better than the conduct the site appears to object to.

More generally, the privacy of the individuals portrayed in the photos is a frequent site of contestation. Many of the screenshots posted to the site include unobscured photos of users’ faces — some of which become the subject of mockery in their own right. As one commenter put it,

²² User comment, posted on 20 July 2011 at <http://www.douchebagsofgrindr.com/contact/>

²³ User comment, posted on 20 July 2011 at <http://www.douchebagsofgrindr.com/contact/>

I found myself enjoying this site but then became uncomfortable at the fact that it is douche-y in it's own kind of tabloid-y, expose-all kind of way. 'Let's laugh at these guys.' It's not really constructive and somewhat hypocritical. ... I'd improve it by putting a black line over the eyes, so that the identities aren't revealed if there are faces (then you can focus on the actual douche-y-ness that is written).²⁴

The commenter identifies a negative stylistic tendency on the site — a mocking approach that detracts from what he understands as the blog's "constructive" goals — and suggests that greater privacy protections for the individuals portrayed would be helpful.

Implementing a black bar over users' eyes — a widely-employed but dubiously effective tactic for protecting identities, including in images published in academic work (Solove, 2008) — would allow visitors to the site to focus on general tropes of bad behavior, rather than on individual instances of misconduct.

A number of the photos also include revealing details that could be used to easily connect a screenshot of a profile to the actual identity of a user offline. One screenshot of a 19-year-old user shows him wearing a sweatshirt with a school's name printed on it, with a barcoded identification badge hanging around his neck. While users themselves chose to share this information on their Grindr profile, its permanent and public display on a blog raises serious questions about whether individuals ought to have the ability to control the circulation of their likenesses. In the words of another commenter,

Although these posts are really funny I seriously believe the persons running this site should at least blur out some of the person's faces. Thankfully I am not on the site, but I can see where this site could go wrong.²⁵

²⁴ User comment, posted on 13 August 2011 at <http://www.douchebagsofgrindr.com/contact/>

²⁵ User comment, posted on 25 July 2011 at <http://www.douchebagsofgrindr.com/contact/>

Manipulating the photos to blur out faces, the commenter suggests, would be a needed step in the direction of protecting the privacy of individuals portrayed on the site. Implicit in the comment is a claim about the ethical stakes for the site: by describing the posts as “really funny,” the commenter sets a fairly high ethical bar for justifying sharing Grindr users’ personal information. The rationalizations of the site’s practices offered by other commenters posit a lower ethical threshold, on the grounds that misbehaving Grindr users are intrinsically less deserving of privacy and protection.

The blog’s actions are also the subject of armchair legal scrutiny. Some readers, for example, dispute whether the site’s moderators in fact have a legal or ethical imperative to obscure personally-identifying information in the screenshots they post on the grounds that Grindr profiles exist as public entities.

The guys can’t sue. Who would they sue, and on what grounds? Can’t happen. They’re posting a profile on an app available to the public, and no privacy should be expected.²⁶

The implicit claim herein is that any interested party could download the Grindr application and browse user profiles; therefore, individuals who choose to share information on Grindr should not have any expectation that their information will remain private or bounded within the sociotechnical confines of the Grindr platform. (Dries Verhoeven used the same argument to justify his display of Grindr profile information in the “Wanna Play?” installation discussed earlier.)

The Grindr privacy policy broadly reflects this understanding of publicity. In describing the types of information gathered, stored, and displayed on Grindr, the service’s privacy policy states that, “When you use the Grindr app, as a default, your

²⁶ User comment, posted on 25 September 2011 at <http://www.douchebagsofgrindr.com/contact/>

profile information is public and other users of the Grindr app can see your profile information” (Grindr, 2013a). Data on Grindr is public by default, and the Grindr profile (as an assemblage of public information) is described as a public entity. This understanding of publicness is also reflected in the “safety tips” offered in the Grindr Help Center:

Most people would not tell a complete stranger their full name, phone number, email address or other sensitive personal information (including bank account details), so be wary when posting this info on your profile. Being careful about revealing your identity is a smart move [sic] when chatting on Grindr. (Grindr, 2014d)

These documents put the burden for managing the circulation of personal information on users themselves. Keeping something decisively private, Grindr’s developers suggest, might require users to keep it off of Grindr altogether. This indicates to users that they should not have an expectation of privacy when they voluntarily share information through the Grindr service.

Some commenters go further, altogether dismissing privacy as an ethical framework by suggesting that users who behave badly on Grindr do not deserve (and therefore should not expect to receive) privacy:

For the morons who believe that these assholes deserve privacy, you’re delusional. Cry me a fucking ocean. They create a public profile parading their racism and superficiality, with which they associate their likeness, and so therefore they deserve every piece of garbage and shit thrown at them. They have a right to do that, and we have a right to scorn them.²⁷

In this case, even if ordinary Grindr users were allowed a reasonable expectation of privacy (something that is not guaranteed in the Grindr privacy policy or terms of service), the commenter suggests that this expectation is negated whenever users publicly

²⁷ User comment, posted on 26 September 2011 at <http://www.douchebagsofgrindr.com/contact/>

post inappropriate content. Further, the commenter insists that the users whose profiles appear on *Douchebags of Grindr* made an active choice to link their inappropriate comments with their likenesses — an act of self-revelation that further undermines any hypothetical claims to privacy they might make.

Ultimately, these contestations leave us with a seemingly intractable ethical morass. “Douchey” misconduct is balanced against even a racist user’s right to data sovereignty; the questionable stylistic approach of the blog is balanced against its contribution to the discursive negotiation of behavioral norms. The durability of these debates over the blog’s three year history suggest that engaging with these questions of data ethics in the specific terms of individual misbehavior on Grindr has not been fruitful.

Taking a step back, we should recognize that these considerations — and the ethical complications associated with them — are not entirely unique to gay male culture. While *Douchebags of Grindr* is narrowly focused on chronicling perceived misconduct among male users on Grindr, similar name-and-shame approaches to counteracting anonymous bad conduct have emerged in a variety of online social spheres. The Tumblr account “Guys of Tinder”²⁸ offers a similarly-structured archive of perceived offensive, inappropriate, or “douchey” profiles, drawn from the mobile dating app Tinder. Mainstream blogs, including *Gawker*, *Jezebel*, and the *Huffington Post* regularly syndicate content drawn from “Guys of Tinder” and other similar blogs (or user’s own submissions). The content featured on these blogs is as unfailingly cringe-worthy as the material on *Douchebags of Grindr*. (One particularly memorable profile encourages a female viewer to “swipe with your lady parts not your heart.”) But, as in the case of

²⁸ <http://tinderguys.tumblr.com>

Douchebags of Grindr, data drawn from Tinder profiles and conversations circulates freely because of perceived wrongs; the author of the blog is judge and jury for the conduct of the users in question, based on contingent and highly subjective ethics.

In place of a continued attempt to solve contingent ethical dilemmas, I advocate for an analytic approach that evaluates these uses of data, and Douchebags of Grindr in particular, in terms of the properties of the inter-platform data flow that enables it. The following section expands upon the architectural properties of these data flows, and outlines an approach to enhance platform boundaries to mitigate the negative impacts on individual data sovereignty that they create.

Solutions: Enhancing platform boundaries

Understanding the risks created by Douchebags of Grindr requires us to disaggregate the differing social functions of the site (normative adjudication and crass entertainment) from the technical processes that enable the display of Grindr users' personal information on the blog. Focusing on the highly variable and subjective social outcomes of the site leads to the ethical quandaries I described in the previous section. By emphasizing the architectural characteristics of data flows between Grindr and the Douchebags of Grindr blog, we can identify which particular informational practices give rise to risk — independent of the social outcomes of those practices. Based on this analysis, I propose two solutions that increase the technical and social burdens associated with moving personal information out of Grindr, while still maintaining needed possibilities for inter-platform data flows to take place. I argue that our overall goal, when faced with inter-

platform data flows, should be to manage them to reduce risk, rather than close them off entirely.

Again, we should begin by precisely mapping out the properties of this inter-platform data flow:

- What types of data are used on the platform?

The Grindr application gathers a wide range of data from users, defined in detail in the service's privacy policy as "Profile Information" (Grindr, 2013a). Generally, this information includes a photo (uploaded from the user's device), textual and numerical descriptions, URLs for other social networking profiles, and automatically-gathered geolocation information.

- How is data gathered, stored, retrieved, and displayed on the platform?

Information on Grindr is gathered through a combination of automatic processes (for instance, the transmission of GPS coordinates from a user's device to the Grindr service) and voluntary user disclosures. Once gathered, a user's profile information is stored privately on the Grindr servers. This server-side storage cannot be accessed directly by users or members of the public; it can only be retrieved by the Grindr client application. When a user opens the Grindr client application, the Grindr app retrieves the profile information of nearby users and synthesizes those data points into the unified profile view available to users.

On Douchebags of Grindr, the unified profile view in the Grindr app is displayed as a JPEG image. These images are stored using the WordPress blogging platform in publicly-accessible and search-engine-indexable folders on the Douchebags of Grindr

web server.²⁹ The images are embedded, using the WordPress platform, into publicly-visible and indexable blog posts.

- What are the steps required to export or remove data from the platform?

Because Grindr's server-side storage cannot be accessed directly by users, the only means of exporting information from Grindr is to capture data directly from the Grindr client application. In practice, this involves using the screen capture function of a mobile device to save a static image (typically in JPEG or PNG format) of the content visible on a device's screen at the moment the screenshot is initiated. The resulting images can be saved on or transmitted from the user's device.

Information on Douchebags of Grindr can easily be copied from the site and redisplayed across other platforms (as was the case when the blog was discussed on other websites). Because profile screenshots posted to the blog are available in a standard JPEG format, they can readily be downloaded and reshared on other platforms. No technical protections exist to prevent visitors to Douchebags of Grindr from downloading or storing content from the site.

- What are the formal rules governing the storage and display of this information over the course of its lifecycle?

Grindr outlines its data storage and display principles in the app's terms of service.

Specifically, the Grindr terms of service prohibit individuals and groups from using

²⁹ The ability for search engine crawlers to index site content is governed by a file called robots.txt, housed at the root level of a website. This file allows a site's operators to block crawlers from indexing specific folders, thereby limiting the automatic spread of a site's content. The Douchebags of Grindr robots.txt file does not contain any restrictions on search engine indexing.

information disclosed on Grindr for any “non-private” purposes. Further, the TOS note that that

User Submissions are owned by the User who submitted them, subject to Grindr’s license to such User Submissions under this Agreement. User Submissions cannot be shared, displayed or duplicated by any other party other than the submitted User. (Grindr, 2014a)

Ostensibly, this clause prohibits exactly the types of third-party display of user information taking place on sites like Douchebags of Grindr. It suggests that users own the information they submit to Grindr, and in using the service grant Grindr a license to store, transmit, display, and use it on their behalf. Users can revoke Grindr’s access to their personal information by deleting their profiles from the service.

Douchebags of Grindr, by contrast, does not offer any indication about how it manages personal information displayed on the site. Specifically, as a number of commenters on the site point out, Douchebags of Grindr does not give users any information about how they could request that their profiles be removed from the site.

The risks in this data flow derive from the changing rules for information management between the Grindr service and the Douchebags of Grindr blog. When someone takes a screenshot of the Grindr application, the protections afforded by the technical structure of the Grindr service — namely, the limited capacity to access Grindr user data exclusively within a Grindr client actively connected to Grindr’s servers — are subverted, resulting in the creation of unrestricted, open-format images, capable of circulating across different platforms with few (if any) barriers to their spread. I want to outline two possibilities for addressing the risks posed by this data flow: (1) A policy-driven solution, which puts the burden for protecting platform boundaries on Grindr

itself; and (2) a sociotechnical solution, which uses software solutions to foster the organic development of user-driven information-sharing norms that enhance the boundaries between platforms while still allowing for a range of user behaviors.

The policy-driven solution relies on Grindr's terms of service to more clearly articulate the rules for data use in the app. As of June 2014, the TOS include three conditions for using the Grindr service that, depending on a reader's interpretation, could prevent potentially harmful inter-platform data flows like the profile screenshots on Douchebags of Grindr:

1. *Personal information ownership*: Grindr asserts that users own the content they submit to Grindr. In using the app, users grant Grindr an "irrevocable, nonexclusive, royalty-free and fully paid worldwide license" to use that content to enable participation in the Grindr service. This does not, however, grant other users who access information on Grindr a license to store, transmit, or reuse that information. Grindr could emphasize this lack of a reciprocal license between users, making it clearer to users that they can reasonably assert ownership over their personal information on Grindr, including when that information is used by other individuals outside of the bounds of the Grindr service.
2. *Application information ownership*: Grindr asserts its ownership over the Grindr software, including the app's interface, design, and information structures. These rules are primarily intended to prevent other developers from releasing copycat applications that directly infringe on Grindr's intellectual property rights. They could, however, also allow Grindr to frame the use of app screenshots on

Douchebags of Grindr as an unauthorized use of the service’s intellectual property (as it extends to the app’s interface and conventions for data display).

3. *Public display*: Grindr prohibits the display of the Grindr application and profile information on any “external display or monitor,” as well as in public settings.

While the clause is designed to prevent public displays like Dries Verhoeven’s “Wanna Play?” installation, Grindr could choose to expand the definition of “public settings” to include the redisplay of user data on blogs and other websites.

In each of these cases, Grindr’s developers could strengthen the wording of these rules, making it clear that sites like Douchebags of Grindr are violations of the service’s Terms of Service.

The problem with each of these approaches is enforcement. In the case of a solution based on personal information ownership, this would require individual users to pursue takedown proceedings directly against sites like Douchebags of Grindr (as well as any sites that repost their personal information). While these procedures have become increasingly common, the burden they place on individual users who do not necessarily have access to the legal resources needed to take these actions is significant. This transforms the management of inter-platform data flows into another task of the “care of the self” — a neoliberal assignment to diligently monitor one’s data and pursue appropriate legal remedies that imposes additional labor on Grindr users. Further, even if users are willing to assume these burdens, the legal status of these processes is still being adjudicated (Babwah, 2010). The ambiguous publicness of self-disclosure on networked services complicates claims that focus on an individual’s lack of consent to being displayed. Remedies that require individuals to directly contest ownership of screenshots

of their profiles run into boundaries of legality that make it difficult — if not impossible — to successfully assert sovereignty over the hybrid products of personal information, commercial software, and individual activity.

The other remedies, focusing on Grindr's articulated policies regarding the ownership and display of information in the Grindr app, have a similarly problematic effect of putting the burden of enforcement onto Grindr itself. In order for either of these approaches to be used against sites like Douchebags of Grindr, the service's developers and operators would be responsible for pursuing legal processes to protect their users. This isn't entirely outside of the realm of possibility, but it assumes that Grindr's developers are willing to spend time and money defending the privacy rights of users who are already violating the service's behavioral guidelines.

A different solution — and one that I would argue holds more promise — uses a technical implementation of boundary-enhancing features to prompt both greater awareness of inter-platform data flows, as well as foster dialogue between users when these flows are initiated. While it isn't possible (on iOS devices, at least) to prevent users from taking screenshots of the currently-running application, it is possible to make this action visible to the users whose data is involved. This approach was pioneered by the messaging application Snapchat, which sends users a notification whenever someone else takes a screenshot of their photos or messages. Implementing a similar notification on Grindr would be a straightforward process: Apple's iOS, for example, has a method built into the operating system's UIKit framework that is specifically designed to notify the currently-running application when a user takes a screenshot. By watching the self-explanatorily-named `UIApplicationUserDidTakeScreenshotNotification` notifier — likely

the same solution implemented in Snapchat’s iOS client — Grindr could automatically recognize any instance in which a user has taken a screenshot of the Grindr application.³⁰ Combined with other readily-available information about the app’s current state — for instance, the ID of the profile currently being viewed — Grindr’s developers could implement a feature which automatically notifies a user whose profile had been screenshotted both that a screenshot has been captured of his profile, and the username of the individual who captured it. Taking this implementation a step further, one could imagine inserting this notification in a chat initiated automatically between the two users, laying the groundwork for a real-time conversation concerning the uses and permissions surrounding that screenshot — as well as, perhaps, the underlying normative objections that prompted the screenshot in the first place.

These notifications do not (and could not) thwart all inter-platform flows. The application sandboxing requirements in Apple’s iOS that are designed to promote privacy and security also keep application developers from stopping users from taking screenshots. Particularly dedicated individuals could use an external camera to take a photo of a device’s display, even in the event that screenshots could be blocked. Nevertheless, these notifications increase the visibility of actions that may cause inter-platform data flows. By informing users that their data may change platforms as a result

³⁰ While the technical details of this implementation go beyond the scope of this project, it’s worth noting that watching a UIKit-wide notifier for user activity does not give an application developers unlimited access to the actions of users. In this case, the UIKit framework passes a notification to the application when it detects that a user has pressed the combination of buttons required to take a screenshot; the application itself does not have the ability to directly monitor the user’s button-presses (or any other system-wide activity), as a result of iOS’s strict sandboxing of third-party applications. More generally, this approach relies on a notification from a device’s operating system that an action correlated with the initiation of inter-platform data flow has taken place, and uses that notification (rather than any direct or ongoing scrutiny of user behavior or information) to trigger an in-app effect. This protects user privacy while still providing developers with the needed information to know when a data flow has been initiated.

of another user's actions, notifications empower individuals to pursue social solutions to unauthorized or unwelcome uses of their information. This transforms the silent "bad manners" of initiating an inter-platform data flow into a visible practice, bringing to light otherwise tacit informational norms and rendering them open for discussion and contestation.

While the remedies I outline in this section could apply to a wide range of data types across many different platforms, I want to emphasize the particularities of this case study. Specifically, I want to suggest that this case is instructive because it illustrates the possibility for inter-platform data flows to also become highly risky inter-contextual flows. When content from Douchebags on Grindr is reposted or linked to on mainstream websites like the Huffington Post, it often does so in the context of entertainment, rather than normative negotiation. Visitors to Douchebags of Grindr who are not Grindr users may have opinions about the appropriateness of the conduct portrayed in the profile screenshots on the site, but they are not, in any meaningful sense, participants in the social context of gay interactivity on Grindr. This calls into questions justifications of these flows that focus on the social functions of the site as a tool for negotiating norms of interpersonal conduct on Grindr; public entertainment, clearly, does not meet this standard.

Crucially, however, the source of risk in these data flows remains the change in the technical rules of data management between the Grindr app and the Douchebags of Grindr blog. The permanent storage of static profile information (screenshots) and the ability to access that information through a public blog increases the possibility for Grindr users' personal information to be viewed or used by unexpected audiences in

unexpected ways. Even if the result of this data flow is a change in both information's platform and context, a granular analytic approach allows us to identify the properties of the initial data flow that create the possibility for escalating risk. This enables us to identify solutions that directly target the actual sources of risk, rather than attempting to engage with the ancillary effects of those risks.

Conclusions: Making sense of revenge porn

The objective of this chapter has been to suggest that, by examining the architectural properties of social contexts and networked platforms, we can identify sources of and remedies for individual risk that are not rooted in the particularities of a given case. Instead of declaring whether a particular use of data is right or wrong, this approach emphasizes the tensions created by the features of the systems that contain it. This allows us to develop generalizable solutions to risky social and technical structures that apply across a range of different platforms, contexts, and cases.

Thus far, the two case studies I've examined — Grindr's business model and the Douchebags of Grindr blog — engage specifically with the ramifications of inter-contextual and inter-platform data flows for gay men's use of networked media. If successful, the model outlined in this chapter, and the solutions offered on that basis, should be able to parse data flows that are not limited to the specific context of gay-targeted social media. Accordingly, I want to end this chapter by taking a step back from the gay-male-focused case studies I've engaged with throughout this work. Specifically, I want to briefly turn to revenge porn, a widely-discussed ethical quandary of unauthorized personal information use, to suggest how a platform- and context-sensitive analytic

approach can aid us in parsing the riskiness of networked data flows, independent of the particularities and complications of a given case study.

Revenge porn represents an especially troubling instance of data failing to stay where its creators put it. The term describes the malicious practice of sharing nude or sexual images of someone that were intended for private use. The typical revenge porn narrative involves a jilted ex-lover posting intimate photos or videos in public fora, causing extensive damage to the reputation of the person portrayed in the photos. Victims of revenge porn often have limited capacities to remove images and videos portraying them once they have been posted online, as the content is frequently shared across a wide range of websites, blogs, and social media accounts. Shame, embarrassment, and professional harm are the seemingly inevitable results of these episodes (Stroud, 2014).

Revenge porn feels particularly troubling because it highlights the ease with the social bonds of intimate trust can be broken. When we share nude photos, we want to trust that the recipient of those photos will recognize that, in most cases, they aren't intended for public display. By sharing these photos, we believe that their exchange takes place within a social context predicated on the privacy and trust of an intimate relationship between individuals. Revenge porn implies fundamental breach of that context.

Focusing solely on the changing social context of sexual photographs obscures the data flows that are actually taking place. As was the case on Douchebags of Grindr, the underlying mechanics of revenge porn center around the transfer of information from one communicative platform (with a limited audience and specific data management principles) to other, considerably more public platforms. The shift from one-to-one

messages to a publicly-accessible blog drastically alters both the potential audience and the information management principles of the content in question. This inter-platform data flow creates the necessary conditions for private images to be reframed as revenge porn.

In order to develop appropriate solutions to revenge porn, we need to focus on the characteristics of the data flows in question, rather than on the emotional gravity of the practice's social outcomes. While a number of states have implemented statutes banning revenge porn — with California securing the first jail sentence for an individual on the basis of a revenge porn conviction in December 2014 — these laws go awry in focusing solely on the harmful social outcomes of revenge porn. A press release from the Office of the City Attorney in Los Angeles frames laws restricting revenge porn as tools to protect reputations from malicious behavior:

“California’s new revenge porn law gives prosecutors a valuable tool to protect victims whose lives and reputations have been upended by a person they once trusted,” said City Attorney Feuer. “This conviction sends a strong message that this type of malicious behavior will not be tolerated.” (Mateljan, 2014)

This approach sets a high bar for the victims of revenge porn. In addition to showing that their personal information has been publicly shared without their consent, individuals need to demonstrate that this was done maliciously by the perpetrator. While malicious intent was fairly easy to demonstrate in the Los Angeles case — the victim had previously secured a restraining order against the defendant, prior to his use of the photographs to shame the plaintiff — it may not always be so straightforward for victims to demonstrate that the perpetrators of revenge porn intended to behave maliciously.

While these statutes are valuable for addressing the harms of revenge porn after they've taken place, a forward-looking approach requires us to examine the informational architecture of these exchanges of sexual images. Using the analytic approach outlined in this chapter, we can describe revenge porn as an inter-platform data flow that, by design, also becomes an inter-contextual flow. Existing statutory approaches focus on the effects of the inter-contextual flow; but, prior to this, we should look for ways to reduce the risks involved in sharing intimate images. This is best achieved by enhancing the boundaries around the communicative platforms individuals use to share this content. One approach would be to implement stronger platform controls in mainstream messaging services — for instance, Apple iMessage, WhatsApp, and Facebook Messenger. While millions of Snapchat users are already taking steps to protect their personal information from unexpected display or storage, the wider implementation of similar solutions could make these protections a communicative default, rather than an affordance of specific applications that individuals need to seek out.

These solutions minimize the risks associated with sharing personal information by approaching the harm of revenge porn from two perspectives. First, platform-enhancing technical solutions maximize possibilities for individual data sovereignty, by giving people greater control over the rules governing the storage and display of their personal information. Second, statutory injunctions on revenge porn give individuals whose personal information has been shared without permission a formal opportunity to pursue legal redress. Together, these solutions address both the inter-platform and inter-contextual elements of revenge porn, without closing off possibilities for individual action.

The overall goal herein is to reduce the need to focus on ethically contingent factors like malice when addressing the risks involved in networked communication. The rightness or wrongness of particular actions in particular cases too often depends on an observer's standpoint; conflicting motivations result in seemingly intractable ethical dilemmas. These dilemmas need not reduce to either absolute relativism, in which any use of data is deemed appropriate, or to presumptions of fault, in which all data flows are assumed to be overly risky and undesirable. By examining risk architecturally, this approach sidesteps heavy-handed ethical determinations and focuses on the on-the-ground characteristics of information flows that give rise to safety concerns. This enables users, policymakers, and software developers to identify narrowly-tailored solutions that directly address actual sources of risk. But, more significantly, by separating the architecture of networked risk from the particularities of individual cases, we can begin to identify generalizable solutions to safety concerns. Instead of continually developing ad-hoc solutions to emerging crises of information use, we can use the typology of inter-platform and inter-context data flows to map new concerns onto existing solutions. Reducing the uncertainties surrounding novel communication practices makes it easier for users and developers alike to protect individual safety in the face of new and unexpected uses of media.

CHAPTER 5

CONCLUSION

On October 16, 2015, Match Group, Inc., the parent company of dating and matchmaking services Match, OkCupid, and Tinder (among 45 other online dating brands), filed S-1 paperwork with the Securities and Exchange Commission, detailing its plan to raise \$100 million through an initial public offering.³¹ As of late 2015, Match Group counted more than 59 million monthly active users and 4.7 million paying users across the different platforms in its network (Statt, 2015). Should the IPO proceed as planned, Match (and its owned brands) will join Facebook, Twitter, and LinkedIn on the relatively short list of social platforms whose business has scaled to the point of warranting its own stock symbol. The ubiquity of these platforms makes it easy to overlook their relative novelty. Match.com, the first Match Group product, launched in 1995. OkCupid launched in 2004, the same year as Facebook. Most recently of all, Grindr went live on the Apple App Store in March 2009. Taking a step back from the discussions over the last four chapters, we need to ask (of ourselves, and of academic research in these fields): In the 20 years since Match.com's founding, and the six years since Grindr's public release, what have we learned about the interplay of identity and information online? And, more critically, what still remains unknown?

This dissertation has examined the interplay of bodies, identities, and digital information as part of a process of networked self-expression on gay social networking applications. At a high level, this project offers a model of networked interactivity that conceptualizes self-expression as an act determined by three sets of affordances and

³¹ <http://www.sec.gov/Archives/edgar/data/1575189/000104746915007908/a2226226zs-1.htm>

constraints: (1) technocommercial structures of software and business; (2) cultural and subcultural norms, mores, histories, and standards of acceptable and expected conduct; and (3) sociopolitical tendencies that appear to be (and are popularly, but inaccurately) understood to be) fixed technocommercial structures. As a way to organize this discussion, the preceding three chapters have used the idea of a lifecycle as a structuring metaphor for describing the creation, collection, management, and use of gay men's personal information. Each chapter engaged with a different component of this process: "Birth" described the relationship among hardware devices, software interfaces, bodies, and identities that is expressed in and through the Grindr app. "Life," in turn, examined the practices of content management that constrain what users are permitted to share about themselves, and in turn described how users creatively engage with those constraints. Finally, "Afterlife" picked up where the previous two chapters left off, examining those uses of gay data that go beyond the manifest functions of gay-targeted social networking services. Together, these discussions offer both theoretical and empirical interventions into more than two decades of scholarship about networked identity, as well as a critical examination of the intersection of digital information and Western gay culture.

Chapter 2, "Birth," engages with a seeming canonic text in studies of online identity: Erving Goffman's *The Presentation of Self in Everyday Life* (1959). Whether explicitly cited or implicitly incorporated, I suggest that Goffman's account of identity as a structured performance looms large in most recent studies of online sociability. "Birth" takes up a series of questions aimed at building a subtler theory of online identity: What can we gain by looking beyond the Goffmanian sociology of performance as a theoretical

framework for conceptualizing identity construction online? How should the methodical study of software interfaces and computational infrastructures figure into an account of the communication of online identity? And, what are the specific complications of subcultural expression in the digital age?

To frame this discussion, I adopt the organic term “birth,” in place of the more rigidly structured “performance.” I argue that the birth of data is a vital process of translation, taking place at the nexus of technical systems, interpersonal dynamics, and personal reflection — where code and humans intersect. Data’s birth is not a totally agential performance; it’s one whose contours are intimately shaped by the interface conventions made available to users through the active decision-making of designers and developers. Users, I suggest, do not merely perform their identities through software interfaces; or at least, they do not perform them in the unified manner of Goffman’s front-stage acting. Instead, people reveal, selectively, under constraint, and through a process of self-surveillance, discrete pieces of information about themselves in accordance with the specific affordances of the software and hardware at their disposal. This account bridges the methodical sociology of Goffman with the Foucauldian notion that the self-revelatory practices of people are intimately shaped by institutional and social contexts. Disclosure, in this context, isn’t merely a performance; it’s a performative act that we’re compelled to do in very particular ways, based on a regime of self-surveillance that shapes not only what we choose to share about ourselves, but also the terms available to us with which we can do that sharing. Critically, I expand on Foucault’s discussion of confession in *The History of Sexuality* (1978b) by suggesting that the material infrastructure of confession — in this case, the hardware and software of

a given platform — plays an essential part in structuring the act of self-disclosure. We are not Goffmanian free agents, not only because our presentations of self are seldom as carefully structured and considered as a theatre performance, but also because the tools at our disposal play a critical part in curtailing agency and creating new structures for creating and sharing individual subjectivity.

This post-Goffmanian theory of online expression gives us the tools to begin to deal with an ever-changing array of platforms, interfaces, and data types. Even since I started work on this dissertation, new platforms have cropped up, offering unique elaborations on embodied expression. The gay-targeted geosocial networking application Grunt, for example, offers users the ability to include a recorded voice introduction on their profiles, integrating vocal expression into the framework of traditionally text-and-image-based profiles.³² A representative for the platform handing out postcards advertising the app at the 2015 Up Your Alley Fair in San Francisco suggested to me that, in the arena of geosocial apps, hearing (not reading, and not seeing) is believing. It's worth noting that profiles on Grunt also include both images and text; so, perhaps, we can more accurately restate the representative's pitch as "hearing, reading, and seeing is believing."

Understanding apps like Grunt — and, more generally, the tendency toward incorporating ever more information into the compact interfaces of mobile applications — might require us to reconsider some of the fundamental assumptions on which most recent scholarship about networked identity and embodiment rest. I began this dissertation by citing Jason Farman's astute argument that both online and face-to-face,

³² <http://www.gruntapp.com>

“full, embodied presence is always being deferred” (Farman, 2012, p. 30). My argument, based on my research into Grindr, is that this position remains true today. But, what happens when the degree of deferral is equal for online and offline interactions? Grunt’s fetishization of voice is a limited first step in this direction; but the combination of video, voice, and other technologies (such as heart rate-sensing wearables) can bring us closer to a state where the limitations of technologies and interfaces do not radically alter the extent to which individuals are able to share and understand information about each other. In that moment, do we discard the “networked” in “networked identity”? Or do we continue to privilege the status of devices in academic research into self-expression and interpersonal interaction? Technological progress, enacted in apps that are available for sale in app stores today, will drive these theoretical questions — and this dissertation takes a critical step in developing the theoretical and empirical tools to engage with them.

At the core of this discussion is a much broader theoretical reframing that I argue is needed to account for the work of the ever-expanding array of social platforms available today. We should begin to conceptualize the deeply intertwined relationship between individuals, groups, devices, and software interfaces as a type of embodied prosthetic: a technosocial extension of the self that is felt and lived as an essential part of individuals, even as its difference and apartness persist. I’m thinking here of what Diane Nelson has called “stumped identities”: those pieces of an individual’s identity that, like a prosthetic limb, “remain lumpy and semi-autonomous,” even as they become a basic part of who we are (Nelson, 2001, p. 319). Our networked presences are unavoidably stumped, their data-driven contours only approximately aligned with the lived

complexities of our bodies and identities. Yet, they remain a part of us — a part that allows us to interact with each other in previously unimaginable ways.

This is not merely an abstract theoretical argument; the stakes for reimagining software as a prosthetic are significant. Like ongoing medical research into building ever-better prosthetic limbs, we should imagine the process of designing and developing social networking sites as a never-completed (and perhaps impossible-to-complete) task of seeking a perfect fit between prosthesis and person. While “Birth” took up a specific set of profile design questions, the set of issues here is nearly limitless. How, and in what forms, embodied data is used to craft better social prosthetics is an open question, and one which scholarship will continue to grapple with in coming years.

Chapter 3, “Life,” argues that the programmatic management of user-generated content plays a central role in structuring practices of networked identity expression. This discussion offers an account of user behavior at its most restricted — and interrogates the ways in which the management of gay data is embedded within broad systems of institutional and commercial power, as well as interpersonal dynamics of resistance and expression. Specifically, this chapter examines the content policies, terms of service, license agreements, and enforcement practices of three gay-targeted social networking services to reveal both the particularities of content management as an emergent phenomenon, as well as the ways in which these policies and practices enter into public discourse around the affordances and constraints of technical platforms. By examining the specific policies of these services alongside the much broader guidelines of mobile software distribution platforms like the Apple App Store, I demonstrate that ecosystem-wide guidance from actors like Apple and Google is insufficient to account for the full

spectrum of restrictions placed on user behavior by services like Grindr. I argue that we can account for the gulf between what Apple condones and what Grindr permits by examining the idealized social and interactive types envisioned by Grindr's developers. This discussion locates content management as a practice that cannot be located outside of the particular sociocultural milieu within which online platforms operate.

This discussion, though rooted in critical academic inquiry into the ideological underpinnings of popular commercial software, has immediate consequences for the millions of individuals who use Grindr, Scruff, Manhunt, and comparable services. The complexities of the techno-legal-commercial ecosystem of apps, app stores, mobile devices, cellular networks, and national and international laws creates a space within which the true justifications for content management can be obscured from users. For example, through a close reading of media coverage of Grindr's executive staff, as well as interviews with the service's CEO, I demonstrate that Grindr's policies are in large part the result of design choices made by a small group of individuals. Despite these active choices, Grindr's developers have resisted taking direct ownership of these policies. The result is an overall atmosphere of confusion, which results in user inaction. I contend that, in a context of ambiguous responsibility, users do not feel empowered to directly question the normative choices that govern the social platforms they use. The adiabaticization of sociotechnical platforms results in collective inaction.

Yet, crucially, a lack of activist outrage should not be construed as total acquiescence by users to the constraints of social platforms. An essential dimension of this study is the recognition that, through tactical engagement with the affordances and constraints of sociotechnical systems, users find ways to meaningfully express

themselves in the face of imposed limitations. From finding workarounds to missing profile fields to using emoji to express prohibited content, users deploy the tools at their disposal to make platforms into livable, lively spaces. Lawrence Lessig's pithy and widely echoed phrase "code is law" may have been more prescient than even Lessig himself anticipated; like laws, the constraints of sociotechnical systems may be bent without breaking. The rigidity of technical systems leaves space for sociotechnical user innovation.

I believe that this, and other studies of what is beginning to be termed "platform politics" (Gillespie, 2010; van Dijck, 2013b), signal the start of a new period of academic and applied inquiry into online interactivity. The assemblage of policies, practices, and principles expressed in content management is the concrete manifestation of what I see as the second epoch of online interactivity: one in which the utopian rhetoric of endless possibility has been replaced (or at least, substantially augmented) by pragmatic, commercial-legal discourses of permission and constraint. Recently, a great deal of academic time and attention has been devoted to parsing the extent to which people understand the particularities of these discourses; for example, whether users understand the click-through license agreements and privacy policies they accept by signing up for different platforms. The conclusions of the overwhelming majority of these studies are that people are ill-informed and ill-equipped to deal with this legalistic new world. Valuable though they are, these studies are only the beginning. What I map out in Chapter 3 is an account of legal and quasi-legal discourse that locates them as but two actants in a broad web of user and developer activity. The meaning of these policies, and the justifications for them, are open to contestation; and, even where their intent is clear,

users may behave in widely divergent or disobedient ways in response. Internet research should not return to the “Wild West” mentality of the late 1980s and 1990s; but nor should it mistake an ever-growing array of formal regulation for the limits of what can and does take place online. Code and law are but two pieces of the puzzle; we should look to the margins of sociotechnical systems — at the behavior of seropositive gay men on Grindr, for instance — for clues about the futures of networked politics. The ethos of “making do” with unwelcoming public spaces is woven throughout the history of Western gay culture; its intersection with online media is a particularly fruitful site to study user behavior in a culture of programmatic constraint.

Finally, Chapter 4, “Afterlife,” examines how both formal and vernacular practices of data use complicate the idea that users are (or, indeed, can be) in control of the circulation of their personal information online. At a high level, I suggest that personal information is frequently used by agents other than the data’s original creator, for purposes that can dramatically differ from the original goals for which it was disclosed — a pattern of distributed use that makes it challenging to promote structures of responsible data use that adequately protect user safety. In this work, I argue that the emerging concept of data sovereignty represents an actionable framework for balancing the interests of the different parties involved in the circulation of digital information: from various constituencies of users to advertisers to software developers. To date, data sovereignty has been the basis of large-scale lawmaking efforts, particularly in the European Union; I suggest that, independent of its function in legislation, it offers a more constructive account of how data is used (and by whom) in online settings than the dominant paradigm of privacy. A focus on privacy, rooted in American constitutional

notions of individual autonomy, does not give us the tools necessary to understand the flow of information through complex, multi-agent systems — of which social networking services are but one example. As part of an overall project of improving user safety and promoting the transparent disclosure of information use online, I map out how information flows through and across both social context and technical platforms to provide a clear, granular framework for understanding the circulation of personal information in complex, distributed sociotechnical systems.

The first case study in Chapter 4, an examination of the dual ad-supported and freemium business models of Grindr, suggests that service providers may not always make it clear to users how differing commercial interests play into the use of their personal information online. Grindr makes extensive use of its users' personal information for the purposes of ad targeting; but the true extent of that use may be obscured from users by the absence of visible advertisements in the paid version of the Grindr app. I argue that Grindr, like other online service providers, has a responsibility to adequately disclose to its users that their information can and will be used to construct profiles for the purposes of demographic targeting in advertisements, independent of some users' choices to pay to hide those advertisements from view.

The second case study, focusing on the blog Douchebags of Grindr, highlights the risks, controversies, and uses of vernacular practices of engagement with user-generated content. The moderators of Douchebags of Grindr publicly post screenshots of Grindr profiles that they deem offensive or inappropriate. The photos, along with associated commentary, are immortalized in blog format, removing a user's own likeness from his individual control. I contend that, for all its crassness, Douchebags of Grindr plays an

important role in the negotiation of norms of behavior on Grindr — and that, more generally, users strategically mobilize different platforms to meet needs that are unfulfilled by the use of one platform alone. Nevertheless, the storage and display of information on the Douchebags of Grindr blog results in personally-identifiable information about Grindr users appearing without consent in a substantially more public format than the manner in which it was initially disclosed on Grindr. As one possible remedy, I outline a technical solution that would more adequately balance the competing interests of different constituencies of users on Grindr by highlighting instances in which data flows across different platforms, while not closing off possibilities for cross-platform vernacular behavior.

Across each of these chapters, this project has been deeply rooted in a desire to make social networking better for gay men — and, by extension, for the myriad constituencies of users who meet and interact with each other every day across a wide range of platforms, both gay-specific and general-purpose. Where possible, I've gone beyond the analysis and critique of the technosocial system of Grindr, and offered concrete prescriptions to software developers and policymakers who are actively working to build the next generation of interactive platforms. While these practices of applied research are common in experimentally-oriented fields like human-computer interaction (HCI), this work applies the techniques of qualitative communication research to identify potential solutions to user concerns and safety problems as they reveal themselves in the public discourses I examine. From interface design elements to code snippets to revisions of the text of an app's Terms of Service, I want these suggestions to put the theoretical and empirical dimensions of this dissertation into conversation with the actual processes

of software design and development that create the systems I'm researching. My hope is that, in balancing theory with practice, this dissertation will advance a model of applied communication research that can more readily put the sophisticated work taking place within the academy into conversation with the user-facing developments taking place outside it. As researchers, I don't believe that we should be content to critique from the sidelines, or reflect on the state of the internet as it is; we should be deeply invested in the ongoing act of creation of the coming generation of technologies.

CHAPTER 6

CODA: DEATH?

In 2009, a Facebook employee published a post on the company’s blog with the title, “Memories of Friends Departed Endure on Facebook.” In the post, the employee reflected on the death of her best friend, a fellow Facebook employee who had passed away three years prior. The employee recalled sharing her grief with her friends and colleagues, and noted that, through grieving, her peers became aware of an open problem facing users of the platform worldwide:

What do we do about his Facebook profile? We had never really thought about this before in such a personal way. Obviously, we wanted to be able to model people's relationships on Facebook, but how do you deal with an interaction with someone who is no longer able to log on? When someone leaves us, they don't leave our memories or our social network. (K. Chen, 2009)

Facebook’s solution was to implement a feature allowing users to “memorialize” the profile of a deceased loved one. A memorialized profile, notes the Facebook Help Center, is “a place for friends and family to gather and share memories after a person has passed away”³³ — a transformation of a user’s networked presence into a sort of digital mausoleum.³⁴ Users have the ability to designate in advance whether they would like their profiles memorialized upon their death, and can assign what the platform calls a “Legacy Contact” to have executive authority over the post-mortem fate of their profile.

³³ <https://www.facebook.com/help/103897939701143>

³⁴ These memorialized profiles have also, in some cases, become sites for online trolling (what Whitney Phillips terms “RIP trolling”). While more restrictive privacy settings might prevent RIP trolls from accessing Facebook memorials, it’s worth recognizing that user behavior doesn’t always share the noble intentions of software designers. For more on the subject of RIP trolling, see: Phillips, W. (2011). LOLing at tragedy: Facebook trolls, memorial pages and resistance to grief online. *First Monday*, 16(12); and Phillips, W. (2015). *This is Why We Can’t Have Nice Things: Mapping the Relationship between Online Trolling and Mainstream Culture*. Cambridge: MIT Press.

And, perhaps most importantly, Facebook notes that memorialized profiles are removed from advertisements on the platform that feature users — a rare break in the service’s relentless monetization of its users’ personal information.

A rather morbid 2012 estimate suggested that, based on Facebook’s growth figures and death estimates from the Centers for Disease Control, 2.89 million Facebook users would die between January 2012 and January 2013 (Lustig, 2012). As Facebook’s user base continues to grow and its existing users continue to get older, one imagines that the stakes for the platform’s treatment of “digital death” will become ever more significant. This fact has not been lost on other online platforms: In the years following Facebook’s implementation of the “memorialization” feature, most other major online service providers have implemented tools and procedures for managing data after a user’s death. And, most telling of all, businesses like Legacy Locker, Securesafe, and Asset Lock have sprung up to help with this process of digital estate planning (Schofield, 2014). In short, we’ve become increasingly aware that, in the age of cheap storage and cloud computing, information has a curious tendency to outlive its creators.

But what does death mean in the context of geosocial networking services? Beyond the emergent digital estate planning industry, what might it mean to conceptualize the death of gay data? Can data on Grindr “die”? What should we make of claims that apps like Grindr have “killed” gay bars or gay culture? And, of course, what about the ever-present possibility of *actual* death — the physical risk associated with encounters between strangers in space that is the dark side of the interpersonal connections at the core of geosocial networks like Grindr? Death, in its various forms,

has hovered at the periphery of the discussions in the last three chapters; now it deserves some more direct attention.

Symbolic death

Perhaps the most persistent claim about apps like Grindr is that, in order for them to live, something else had to die. Grindr has been accused of killing gay bars, gay urban culture, gay sociability beyond hooking up, and even pre-digital forms of cruising (Kapp, 2011; Thomas, 2011; Trebay, 2014; Vernon, 2010; Yiannopoulos, 2012) — a lengthy list of accusations that bring new meaning to the term “killer app.” “I don’t even bother going out anymore,” a 28-year-old Londoner was quoted in one article as saying about the impact Grindr has had on his social life (Yiannopoulos, 2012) — a suggestion that apps like Grindr obviate the need to participate in the traditional “offline” sites of gay culture in order to pick up guys. The article then contends that Grindr’s deleterious impact is most pronounced in cities and neighborhoods with well-established gay scenes. While apps like Grindr offer users in rural or remote areas (where gay community might be hard to come by) a way to connect with other men, they do so with the collateral cost of undermining those urban communities which already exist.

These issues are consistently framed as questions of how much the gay community is giving up by embracing the convenience of meeting on Grindr. To me, it seems that arguments about the death of gay bars indulge in a sort of generational nostalgia about what gay urban life looked like before the rise of smartphones and Grindr-facilitated hookups. But this framing largely misses the point: Grindr is one part of an elaborate network of tools — some technical, some spatial, some social — that gay

men use to pursue their goals, be they romantic, sexual, social, professional, or otherwise. We need to look at gay bars and neighborhoods without the frame of nostalgia, and instead as part of an evolving set of attitudes and behaviors around the operation of gay sexuality in changing sociopolitical climates of acceptance and normalcy. It's true that some Grindr users might choose to log onto the app instead of going out to a bar; but this speaks more to changing perspectives on the value of sites of local gay community than it does about the impact of new technologies like Grindr. Whether gay bars live or die, it's unclear that their fate has much at all to do with the rise of geosocial networking services — other than that these apps make for convenient scapegoats. In my estimation, the life of gay apps and the death of gay bars are two distinct phenomena.

Within the realm of networked media, however, death takes on a more personal form for users, as a type of social invisibility — an absence from the hustle and bustle of internet sociability that, on Grindr, takes the form of a swift and total expulsion from the app's Cascade of profiles. Grindr's focus on immediacy is unrelenting. The app is built from the ground up to facilitate quick connections — both online and offline — between users. Profiles are relatively terse, comprised primarily of images, quantitative measurements, and character-limited textual descriptions. Users of the paid "Grindr Xtra" service can even save and reuse phrases, allowing them to quickly send the same text to multiple users with minimal effort. In the event that users decide to meet face-to-face, the app offers both a simple way to share one's location on a map, and an estimate of how many minutes it would take someone to walk to the designated location. And, since the free version of the app only displays the 100 nearest profiles, users in urban areas frequently encounter Grindr as a rapidly-shifting, unstable collection of different users

moving in and out of their vicinity. The rapid pace of sociability on Grindr is both exhilarating and exhausting.

As part of this designed emphasis on immediacy, Grindr leaves little room for users who only engage with the app occasionally. Unlike many social networking sites, Grindr removes profiles from public view after 24 hours of inactivity; failing to open the Grindr app at least once a day entirely eliminates one's presence on the Grindr network (at least until the next time the app is launched). Grindr is built to be quick and easy to use; but it's also best used when one goes all-in, engaging with the app at frequent intervals throughout the day, in different locations. It's unsurprising, therefore, that the service's fact sheet suggests that the average Grindr user logs onto the app more than 8 times per day, spending more than 2 hours chatting and browsing profiles (Grindr, 2013b). Failure to engage frequently leads to a form of social death: a diminished visibility that decreases one's chances of seeing, being seen by, and connecting with other users. The technical design of individual presence in the Grindr Cascade is designed to thwart what Ben Light (2014) calls "disconnective practices": everyday tactics of use and non-use of networked media that allow people to engage with services on their own terms. Grindr's design strongly disincentivizes these practices.

This notion of death-by-invisibility is a curious artifact of networked media. Facebook's EdgeRank algorithm rewards users who use the service "well" with heightened visibility; less successful users are punished with a diminished presence in the News Feeds of their friends and followers (Bucher, 2012). Google's PageRank algorithm applies this same ruthless logic to its organization of search results: "relevance," assessed in a variety of ways (including keywords, incoming and outgoing links, publication date,

and, most recently, whether the website offers a mobile-optimized version of itself), can push a particular page into the coveted first 10 results, or relegate it to a lower position in the heap (Halavais, 2009). Various measures of successful networked presence — follower counts, Klout scores, and so on — have all become a part of an overall social ecosystem in which invisibility is a constant threat looming over the heads of users who fail to invest the proper labor in maintaining their presences on various platforms (Marwick, 2013). It's perhaps unsurprising, therefore, that so many Grindr users spend so many hours on the app — and, relatedly, that one of the chief complaints of users who left the service is that constantly engaging with the app became, for them, a chore and an immense waste of time (Brubaker, Ananny, & Crawford, 2014). Remaining visible demands a lot of work.

For some users, this burden of visibility results in an ambivalent attitude towards gay social networking services. Elija Cassidy (2015; 2013) uses the term “participatory reluctance” to describe the seemingly paradoxical feelings gay men report about these networks. On one hand, users claim to be consistently dissatisfied with what they get from services like Grindr; they insist that, despite the presence of scores of other users, they can never quite find the right someone (or someones). On the other hand, users overwhelmingly tend to stay on gay social networks. Perhaps their vocal dissatisfaction is just idle kvetching; but if users' behavior is at all an indicator of their feelings, the failures of gay social networks never quite outweigh the benefits of staying put. Being altogether absent from the Grindr Cascade would be a form of social death many gay men are unwilling to accept.

As an alternative to leaving a particular service in favor of another, many individuals maintain presences across multiple networks as a way to mitigate the specific shortcomings of different platforms. For example, in interviews with gay men, David Gudelunas (2012b) found that many of his respondents reported maintaining multiple profiles across different social networking services — with some managing as many as ten profiles at once. These profiles, Gudelunas writes, “were not seen as discrete entities, but rather as part of an elaborate network” (2012b, p. 13). In this sense, an individual’s participation in gay networked community isn’t something that can be reduced to a single profile on one website or app; presence is distributed across multiple platforms. The birth or death of any one profile becomes a small part of a much larger process of networked identity formation.

This practice has important consequences for the operation of identity in and across different platforms. Chapter 4 outlined the ways in which the boundaries around particular services are more porous than they might seem by mapping out how flows of data between different platforms spread personal information. But beyond the meta-commentary offered by blogs like Douchebags of Grindr, users also carry information about the people they encounter from service to service. Within a web of profiles on different gay social networks, Gudelunas notes that individuals reported being able to “triangulate” the identities of other users, taking advantage of the particular features of different platforms to construct maps of local gay social networks that extended across individual websites or apps. Rather than focusing on individual profiles, users are able to work across platforms to generate, for themselves, robust inter-platform understandings of networked identity. Data, in this sense, is not a fixed entity that can live or die on a

particular platform, but is part of an elaborate system of sociotechnical relationships that extend across multiple devices, services, and social contexts.

Some users also move between platforms as the character of particular apps changes over time. While Grindr was the first gay-targeted geosocial network, a number of competitors, including Scruff, quickly emerged as alternatives for users seeking more niche communities. Scruff, for example, was initially billed as a network targeted at older, larger, and more hirsute men — a nod toward the well-established bear subculture (Roth, 2014). Initially, Scruff and Grindr attracted at least somewhat different audiences: the grids of profiles on the two apps overlapped in some small part, but users approached the differing marketing rhetoric of the apps in the same way they might approach bars targeting different gay subcultures. And, indeed, the two apps make at least cursory attempts to differentiate themselves beyond simply marketing: Scruff, for example, includes quantity of body hair as a field in user profiles, signaling through the inclusion of this interface element that the intended Scruff user is someone who sees his body hair as something worth disclosing in a terse social network profile — and who, in turn, is interested in that information about the people he meets. Ultimately, these apps do not represent different paradigms within the domain of gay geosocial networking; but through particular design and marketing choices, service providers are nevertheless able to signal to users who the intended audience for an app is.

This signaling isn't always successful. A Scruff user I interviewed in 2014 noted that, while he found users who are “more like [him]” on the app when he first joined it in 2012, the in-app community has increasingly tended to overlap significantly with Grindr. He also suggested that many of the people he recalled seeing on Grindr eventually

migrated onto Scruff. The boundaries between the two services are porous, and only loosely defined by service providers' visions for what their intended user bases look like. This account, which parallels my own sense of how the apps have evolved, offers an important insight into the life and death of gay social networking services. The popularity of different services ebbs and flows, for reasons that aren't always easy to identify. While a relatively small number of players dominate the field of gay social networking, the meteoric rise of Grindr suggests that a new service with innovative features can easily disrupt the entrenched players in the industry. Users move between services, sometimes following their networks, other times in pursuit of entirely new networks of potential connections. The consistent narrative across multiple generations of gay social networks (and, indeed, social networks more generally) has been that users invest significant time and effort in maintaining their profiles — but can and will abandon them for the greener pastures of a new platform should its affordances prove sufficiently compelling.

But what about those users who choose to leave altogether? In interviews conducted with former Grindr users, Brubaker, Ananny, and Crawford (2014) noted that many gave up on the service after concluding that the potential of the in-app community — the promise of meeting someone new for a date or a hook-up — failed to come to fruition in their experience. Users expressed a wide range of complaints about the service: from the overwhelming focus (amongst Grindr users) on casual sex to the significant time demands imposed by the service's social model to the “dehumanizing” and “flesh-focused” structure of the app's profiles. Brubaker, Ananny, and Crawford rightfully recognize these points of dissatisfaction as a mix of technical and social factors: some, like the structure of profiles (or dissatisfaction with the service's content

management policies, discussed in chapter 3), are a result of specific design choices; others, like the preponderance of individuals seeking casual sex, are products of how users choose to employ the ambivalent affordances of the medium of geosocial networking apps. In the end, the authors suggest that users choose to leave Grindr because the actual experience of the app fails to line up with users' idealized visions of what they want gay social networking to be.

Even though Grindr makes users' profiles invisible after only 24 hours, leaving the service altogether is considerably more challenging. Many of the participants in Brubaker, Ananny, and Crawford's study reported that they looked for a way to delete their profiles, but ultimately settled for simply deleting the Grindr app from their phones — an approach that removes a user's point of contact with the Grindr service without removing any of that user's personal information or conversations from Grindr's servers. Traces of personal information — photos, bodily “stats,” records of conversations, and lists of favorite users — persist, even when a user has made the choice to end their participation on Grindr. This information can remain active on Grindr's servers for as long as the service chooses to retain them. In practice, the service has little incentive to delete inactive user accounts; cheap networked storage reduces the need to purge inactive data, and returning users might enjoy the frictionless experience of rejoining the app to find their profiles already populated with all the relevant personal information. Leaving Grindr is certainly possible; but, short of a concerted dig through the app's settings page to find the relevant link to delete one's profile, users' data lives on long after their departure.

The ambiguous fate of personal information after a person's departure from an online platform is by no means unique to gay-targeted social networks. Facebook, for example, makes a concerted effort to persuade users to deactivate their profiles, rather than deleting them altogether. First, unlike all other account-related settings, the service hides the account deletion function in a Help Center page that describes what deleting an account entails. Even savvy users who are accustomed to perusing social network settings might encounter some difficulty finding out how to delete their accounts. The service instead offers users the possibility of deactivation, a temporary suspension of a user's account that keeps that user's existing content and connections in place, waiting for reactivation. It's telling, perhaps, that users who choose to deactivate their profiles are presented with the option of auto-reactivation after a designated period of time. Leaving Facebook altogether is framed as a nuclear option — far too extreme for users who are merely dissatisfied with some element of their experience on the site. Appropriately, these practices of deletion are commonly termed “social media suicide” (Light, 2014). And, even in the case that a user chooses to pull the trigger on their profile, Facebook maintains deleted accounts for 14 days in case users have a change of heart, noting that some information may persist for as long as 90 days after deletion while all traces of a user's presence are wiped from the service. The death of a Facebook profile is a protracted process.

To their credit, a number of major online platforms offer users the ability to export a version of their personal information prior to deleting their account. (Grindr does not provide any formal way for users to export information they stored on or shared through the service.) Google, for example, gives users access to the Takeout tool as a way

to download some or all of their account information. Data portability is an important step in the direction of network openness: using Google Takeout, users can elect to move the information they've shared with Google over the years — whether in the form of search histories, Gmail archives, Google+ posts, Picasa photos, or reviews in Google Maps, to name just a few possibilities — to a competing service. This empowers users to make active, ongoing choices about the online services they choose to use, rather than forcing them to remain on certain platforms by default.

Despite the ease with which tools like Google Takeout give users access to their information, it's worth questioning whether the availability of data is equivalent to actually giving people the ability to make use of it. An archive of a social network account as a collection of XML and JSON objects is a far cry from the integrated experience of that data as a live profile. Platforms perform a great deal of integrative work, translating disparate pieces of data into usable, meaningful entities. Static, machine-readable archives are their own form of data death: the demise of information's utility, in the face of proprietary or constantly-evolving technological standards for information storage or display.

Real death

Thus far, I've treated death as a metaphor: a way to think about the rise and fall of businesses, or how information travels through networked systems. But a very real, non-symbolic form of death — the possibility of physical harm resulting from an encounter that originated online — lurks beneath the surface of every interaction on Grindr. Julian Dibbell's account in "A Rape In Cyberspace" (1994), discussed at the start of this work,

probes the psychosocial consequences of interpersonal violence in a text-based online environment; what happens when that violence migrates from networked settings into the “real world”?

The reported instances of Grindr-enabled assault are harrowing for people familiar with the platform. A string of knifepoint robberies in Sydney, Australia were linked to a man who used Grindr to find and meet his victims (L. Hall, 2015). A Grindr user in London was robbed in his own home by a man wielding a cattle prod after he invited over someone he met on the app (Gremore, 2014). A man in Seattle was beaten with a hammer by someone he met on Grindr — with police noting that the attacker proceeded to viciously bite the victim “when the hammer was not enough” (Pulkkinen, 2014). (Enough *for what?* But the police offered no further elaboration.) A Canadian tourist visiting Philadelphia was raped, beaten, and robbed after inviting a man he met on Grindr to his hotel room (Lattanzio, 2014). In each case, the interactions leading up to the assault were entirely ordinary for Grindr: conversation, an exchange of photos, followed by an arrangement to meet. These episodes are scary precisely because of their banality: up until the moment when you’re being robbed at gunpoint, there are few concrete signals that the connection you just made on Grindr might not be safe. Even cautious users following all of the service’s safety tips (Grindr, 2014d) can’t guarantee that the person they’re chatting with doesn’t have malicious intentions.

The safety risks posed by malicious uses of Grindr become especially acute when gay sociability intersects with institutional, governmental, or cultural contexts that are less than welcoming to gay community. Grindr proudly touts the fact that the service’s users reside in 192 countries around the world — a list that includes seemingly

inhospitable locales like Russia and Iran. But, in practice, this global reach is as much a source of danger as it is a cause for celebration. In Russia, for example, anti-gay vigilante groups like Occupy Pedophilia³⁵ use Grindr and other gay social networking services to lure gay men into meeting, after which they verbally and physically abuse their victims (May, 2015). The groups take advantage of Russia's overtly antagonistic legal position towards LGBT individuals by making no secret of their activities; videos of the beatings are often posted online.³⁶ Herein, Grindr becomes a way for interested organizations — be they groups of private individuals or, perhaps, governments themselves — to identify the locations of gay men. Sharif Mowlabocus (2014b) has argued that the service's cavalier attitude about sharing location information for users in potentially dangerous sociopolitical climates reflects a “privileged, white, middle-class Anglo American” model of gay sociability — an incarnation of Joseph Massad's (2002) “Gay International” in the sphere of digital media. Yet, it's unclear how Grindr's developers should react to charges that they haven't done enough to protect their users when users continue to enthusiastically share their personal information despite a growing awareness of the associated risks.

Questions of physical safety become even more urgent when they're coupled with what Wendy Chun (2006) has called one of the most enduring “paranoid narratives” of digital media: the need to protect minors from the unregulated sexual wilds of the internet. The number of publicized accusations of sexual assault involving Grindr users under the age of 18 is relatively small: only three have resulted in legal proceedings since

³⁵ The group uses the terms “pedophilia” and “homosexuality” interchangeably.

³⁶ https://www.youtube.com/watch?v=zMTbFSJ_Tr4

the service's launch.³⁷ Yet, even in their low numbers, the cases are chilling. In June 2012, two adult Grindr users organized a threesome with another user who claimed in his profile to be over 18, but in fact was under age; both men faced charges of sexual assault and endangering the welfare of a child (Goldman, 2015). In September 2014, a seropositive man was charged with sexual assault after admitting to having sex with a 15-year-old on at least four occasions — including one instance of unprotected sex without having disclosed his serostatus (Molinet, 2014). In April 2015, another man was charged with sexual abuse of a child after he met and sexually assaulted a 14-year-old he met on Grindr (“Social Media App May Have Played Part in Alleged Sexual Assault,” 2015).³⁸

In each case, public accounts of the assaults are framed in largely the same way: older, predatory men take advantage of impressionable, naive teenagers who have stumbled their way onto Grindr. The Grindr app becomes a hotbed for this type of predation, and both the individuals involved and the service itself are held to task. The District Attorney responsible for prosecuting the 2015 case suggested that such incidents were bound to become more common as social networking services like Grindr continue to rise in popularity; the possibility for abuse, he argued, is “the drawback of modern technology.” The onus, therefore, is on parents to teach their children to avoid these kinds of dangerous situations; in the DA’s words, “You know you tell them stay away from strangers on the street. Stay away from strangers on the phone” (“Social Media App May Have Played Part in Alleged Sexual Assault,” 2015).

³⁷ This number reflects only those instances of Grindr users becoming sexually involved with minors in which civil or criminal proceedings were initiated against the adult user. One can assume that a greater — potentially, much greater — number of instances of sexual assault and statutory rape go unreported.

³⁸ Because minors are unable to legally consent to sex, any sexual activity involving a minor is definitionally considered a form of sexual assault.

These accounts echo many of the classic tropes of online child safety narratives: the essentially dangerous nature of new media; the need to impose strict, top-down controls on how minors use the internet; a digital reincarnation of “stranger danger” in the figure of the older male sexual predator; and the importance of raising children to be safety-savvy and highly private. Yet, absent from these discussions is even a cursory recognition that the new medium of gay-targeted social networking may be a crucial social outlet for gay, bisexual, and questioning youth. While gay youth-oriented chat rooms and social networking services were available in the early 2000s, these services have largely fallen by the wayside, in favor of general-purpose platforms like Twitter, Facebook, and Snapchat. Perhaps this is truly representative of an increasingly absent demand among young adults for networked spaces to engage with peers about their sexuality; but it’s worth considering how, if at all, the current generation of popular sites of gay networked sociability might fit into an overall queer social landscape that increasingly includes individuals under the age of 18. Even with the service’s extensive content management, Grindr may well be too lewd or too hook-up-oriented to be a safe and age-appropriate resource for teenagers; but the fact that people under 18 are on these services already indicates that we can’t readily dismiss these platforms out of hand as loci for queer youth culture. Rather than merely trying to absolve themselves of legal responsibility or, worse, trying to drive out teenagers entirely, service providers should instead focus on crafting safety strategies that can accommodate a wide variety of use cases for platforms like Grindr — including, possibly, their role in safely connecting queer young adults.

In its public response to these cases — and particularly in instances involving underage users, Grindr has stressed the importance of its policies (including its content management strategy) as part of an attempt to keep its users safe.

The safety of Grindr's users is paramount and we have a strict terms of use policy that require users to be aged 18 years or older. We have deployed a large team of moderators focused on monitoring and ensuring users adhere to our terms of service guidelines. As an added protection, we encourage parents to add parental controls to their children's devices to help ensure that their children cannot access high maturity sites and apps. We also have a number safeguards in place, including our strict photo and profile content guidelines, privacy policy, detailed user agreements, terms of service and safety tips which can be found on our website. We encourage our members to utilize the safety mechanisms we've made available. (“Social Media App May Have Played Part in Alleged Sexual Assault,” 2015)

Grindr claims that it makes a reasonable effort to screen out underage users, both through active moderation and software implementation of so-called “age-gate” technologies (such as “17+” App Store parental control ratings). Beyond this, the service claims that users are responsible for their own actions, and that parents are responsible for the actions of their children. But, as a rule, the service avoids making any explicit claims of responsibility for the well-being of their users. Like other online service providers, Grindr enjoys the protection of Section 230 of the Communications Decency Act, which provides immunity for service providers from prosecution on the basis of the potentially illegal actions of their users (Goldman, 2015).

It’s worth pausing herein to recognize that the very notion of a “Grindr-enabled assault” is a problematic one. Is Grindr responsible for the illegal or violent actions of its users? Under Section 230 of the Communications Decency Act, no. But can we meaningfully attribute any kind of causal responsibility to the service on the basis that the

individuals involved in each of these cases met on Grindr? Is there a special kind of danger associated with “strangers on the phone”? Put counterfactually: Would these assaults have happened anyway (perhaps with different victims) if Grindr didn’t exist? Assaults, robberies, and rapes happen when people meet offline, too; yet in these cases, we’re fixated on the fact that a new technology seemed to play a crucial part in enabling something tragic to happen. A more specific analytic frame is needed: namely, what is it that’s *different* about the way risky social interactions work online? What are the unique risk-bearing affordances of novel technologies?

The issue at the heart of this discussion, and one that brings together the various strands of the previous three chapters, is how users construct their identities on and through geosocial networking apps. Throughout this dissertation, I’ve examined the processes and practices that govern self-expression and user behavior; but it’s worth making explicit the ways in which the expressive affordances of these platforms can stand in the way of individual safety. A chief source of uncertainty in risk in these encounters is anonymity. While Grindr encourages users to disclose a great deal of personal information — including a photo and their location — none of this information necessarily needs to be personally *identifying*, or, in fact, accurate at all. Oftentimes, these concerns are framed as issues of aesthetic deception; as Lauren Sessions (2009) put it, users who manage their presences too effectively are subject to the charge that they “looked better on MySpace.” But false or incomplete profiles can be dangerous as well as deceptive: malicious users can employ misleading profiles to lure people into face-to-face encounters without the risk that their profile can easily be traced back to their offline identity. Or, in the case of underage users, a misleading profile puts both parties at

significant legal risk. As the episodes discussed above show, the risks of a bad face-to-face encounter go far beyond the possibility that your hook-up is three inches shorter than he claimed.

Other platforms that are built around connecting users in physical space employ a variety of identity credentialing techniques. Tinder, for example, constructs user profiles using information drawn from the Facebook API — an approach that presumes a greater degree of authenticity or truth in what a user posts on Facebook. (This, of course, need not actually be the case; but constructing a convincing false Facebook with a robust network of connections requires enough effort that deception becomes too labor-intensive for anyone but the most committed individuals.) Airbnb, a lodging rental service, relies on identity documents like drivers' licenses to confirm the identities of potential hosts and guests. The ride-sharing platform Uber likewise employs a variety of techniques for confirming the identities of drivers — from government identity documents to criminal background checks to social network searches. None of these approaches are perfect guarantors of safe encounters between people; but, reasonably, they put faith in external signals of identity and good behavior as a way to ensure safety when digitally-mediated interactions between users move offline.

It's not altogether clear if or how Grindr should approach the task of credentialing its users. One of the app's core user experience goals is a frictionless profile creation process: users should be able to establish their presence on the app with a minimal amount of effort and with a high degree of control over how much (or how little) they choose to share on the service (see chapter 2). While some identity verification

techniques, such as requiring users to share their phone number³⁹, are relatively low-impact in terms of the effort they require, any requirement for additional disclosure of personal information undermines the ability of users to choose to remain semi-anonymous. For some subset of Grindr users — openly gay men in welcoming sociopolitical settings with few qualms about online privacy — these changes might be unoffensive; but for many others, any connection between their presence on Grindr and their other online and offline identities might constitute too great a burden. The possibility for limited disclosure is its own affordance, and service providers should tread carefully when considering whether the potential benefits of identity credentialing outweigh the loss of individual sovereignty over how much (or how little) to share.

Grindr was neither the first gay social network, nor the first use of mobile location data to build social ties. But when it launched in 2009, Grindr nevertheless offered its users something truly new: an intuitive, elegant, and — in the words of many users — addictive package of different technologies that gave the app’s users the ability to quickly and easily connect with the people around them. Through the right combination of novel technology, effective design, and suggestive marketing, the service was able to transform itself into the “killer app” of gay networked media. The continuing presence of millions of engaged, active Grindr users worldwide stands as a testament to the potency of this combination.

³⁹ This approach has been implemented by Twitter as a way to curtail the serial creation of so-called “burner” accounts for abusive purposes.

As Grindr matures from a scrappy upstart into an established platform, the problems it and its users face are bound to become more complex, and inevitably will entail a series of trade-offs in which, seemingly, there's no clear "right" answer. This dissertation has explored three such trade-offs: A simple profile design makes participation easy; but adding complexity makes Grindr more accessible to people whose identities defy reductive expression. Content management maintains Grindr's normative commitments to polysemy, making the app more than just an online sex club; but the service's conservative approach constrains the ability of individuals to share a wide range of information, including but by no means limited to nude photos. An open, permissive sociotechnical space on Grindr encourages people to employ technologies in innovative ways to form connections and communities; but users don't always have the tools they need to make informed choices as they navigate an increasingly complicated web of social contexts and technical platforms. The solutions I've suggested here are one possible point of entrance into an ongoing process of building better social networking services. Our more general commitment, as designers, developers, policymakers, critics, and advocates should be to strive to craft sociotechnical systems that respect the users whose data lives and dies on the platforms we create.

WORKS CITED

- Adey, P. (2010). *Aerial life: Spaces, mobilities, affects*. Chichester: Wiley-Blackwell.
- Albrechtslund, A. (2008). Online social networking as participatory surveillance. *First Monday*, 13(3).
- Ambrose, M. L. (2013). It's About Time: Privacy, Information Life Cycles, and the Right to be Forgotten. *Stanford Technology Law Review*, 16.
- Andrejevic, M. (2005). The work of watching one another: Lateral surveillance, risk, and governance. *Surveillance & Society*, 2(4).
- Andrejevic, M. (2011). The work that affective economics does. *Cultural Studies*, 25(4-5). <http://doi.org/10.1080/09502386.2011.600551>
- Anton, A. I., Earp, J. B., He, Q., Stufflebeam, W., Bolchini, D., & Jensen, C. (2004). Financial privacy policies and the need for standardization. *IEEE Security & Privacy*, 2(2).
- Apple. (2013a). App Store Review Guidelines. Retrieved August 19, 2013, from <https://developer.apple.com/appstore/resources/approval/guidelines.html>
- Apple. (2013b). Apple Special Event: October 22, 2013. Retrieved from <http://www.apple.com/apple-events/october-2013/>
- Babwah, L. A. (2010). Climbing in Our Windows & Snatching Our Likenesses up: Viral Videos & the Scope of the Right of Publicity on the Internet. *North Carolina Journal of Law & Technology*, 12, 57–76.
- Barbour, K., & Marshall, D. (2012). The academic online: Constructing persona through the World Wide Web. *First Monday*, 19(9).
- Barry, A. (2001). *Political machines: Governing a technological society*. London: 254

- Athlone Press.
- Bauman, Z. (1995). *Life in Fragments: Essays in Postmodern Modernity*. Oxford: Blackwell.
- Bauman, Z., & Lyon, D. (2013). *Liquid Surveillance*. Cambridge: Polity.
- Baym, N. (2010). *Personal Connections in the Digital Age*. Cambridge: Polity.
- Baym, N. (2011). Social networks 2.0. In M. Consalvo & C. Ess (Eds.), *The handbook of internet studies*. Hoboken.
- Beck, U. (1992). *Risk Society*. London: SAGE.
- Bengry, J. (2009). Courting the Pink Pound: Men Only and the Queer Consumer, 1935-39. *History Workshop Journal*, 68(1), 122–148. <http://doi.org/10.1093/hwj/dbp006>
- Berlant, L., & Warner, M. (1998). Sex in public. *Critical Inquiry*, 24(2).
- Berry, D. M. (2011). *The Philosophy of Software*. New York: Palgrave Macmillan.
- Best, K., & Tozer, N. (2013). Scaling digital walls: Everyday practices of consent and adaptation to digital architectural control. *International Journal of Cultural Studies*, 16(4). <http://doi.org/10.1177/1367877912460618>
- Birnholtz, J., Fitzpatrick, C., Handel, M., & Brubaker, J. R. (2014). Identity, identification and identifiability: The language of self-presentation on a location-based mobile dating app. *MobileHCI 2014*, 3–12. <http://doi.org/10.1145/2628363.2628406>
- Blackwell, C., Birnholtz, J., & Abbott, C. (2014). Seeing and being seen: Co-situation and impression formation using Grindr, a location-aware gay dating app. *New Media and Society*, 17(7), 1117–1136. <http://doi.org/10.1177/1461444814521595>
- Boellstorff, T. (2008). *Coming of Age in Second Life*. Princeton: Princeton University

- Press.
- Boellstorff, T., Nardi, B., Pearce, C., & Taylor, T. L. (2012). *Ethnography and Virtual Worlds*. Princeton: Princeton University Press.
- Boland, R. (2002). Information system use as a hermeneutic process. In M. D. Myers & D. Avison (Eds.), *Qualitative research in information systems A reader*. Thousand Oaks: SAGE.
- Bourdieu, P. (1984). *Distinction*. Cambridge: Harvard University Press.
- boyd, D. (2014a). *It's Complicated*. New Haven: Yale University Press.
- boyd, D. (2007, June 24). Viewing American class divisions through Facebook and MySpace. Retrieved March 19, 2013, from <http://www.danah.org/papers/essays/ClassDivisions.html>
- boyd, D. (2014b, March 21). Why Snapchat is valuable: It's all about attention. Retrieved November 26, 2014, from <http://www.zephoria.org/thoughts/archives/2014/03/21/snapchat-attention.html>
- boyd, D., & Hargittai, E. (2010). Facebook privacy settings: Who cares? *First Monday*, 15(8).
- Browne, S. (2010). Digital Epidermalization: Race, Identity and Biometrics. *Critical Sociology*, 36(1), 131–150. <http://doi.org/10.1177/0896920509347144>
- Brubaker, J. R., Ananny, M., & Crawford, K. (2014). Departing glances: A sociotechnical account of “leaving” Grindr. *New Media and Society*, 1–18. <http://doi.org/10.1177/1461444814542311>
- Bucher, T. (2012). Want to be on the top? Algorithmic power and the threat of invisibility on Facebook. *New Media and Society*, 14(7), 1164–1180.

<http://doi.org/10.1177/1461444812440159>

Bullock, P. (2011, August 27). Puerto Rican Republican has a perfectly good reason for being on Grindr. Retrieved November 26, 2014, from

<http://www.towleroad.com/2011/08/puerto-rican-republican-has-a-perfectly-good-reason-for-being-on-grindr.html>

Burnett, G., Whetstone, M., & Jaeger, P. T. (2013). Personal health record interfaces: A hermeneutic analysis. *First Monday, 18*(8).

Cain, S. (2014, October 9). Grindr, dick pics and contemporary art's new invasions of privacy. Retrieved November 26, 2014, from

<http://www.theguardian.com/artanddesign/2014/oct/09/grindr-dick-pics-privacy-invasion-contemporary-art-dries-verhoeven>

Caldwell, J. T. (2008). *Production Culture: Industrial Reflexivity and Critical Practice in Film and Television*. Durham: Duke University Press.

Calhoun, A. (2008, December 31). Facebook's War on Nipples. *Time*. Retrieved from <http://www.time.com/time/printout/0,8816,1869128,00.html>

Campbell, J. E. (2004). *Getting it on online*. New York: Harrington Park Press.

Capino, J. B. (2005). Homologies of Space: Text and Spectatorship in All-Male Adult Theaters. *Cinema Journal, 45*(1), 50–65.

Cartwright, L. (1995). *Screening the body: Tracing medicine's visual culture*. Minneapolis: University of Minnesota Press.

Cassidy, E. (2015). Social networking sites and participatory reluctance: A case study of Gaydar, user resistance and interface rejection. *New Media and Society, 1*–16.

<http://doi.org/10.1177/1461444815590341>

- Cassidy, E. (2013, August 4). *Gay men, social media and self-presentation: Managing identities in Gaydar, Facebook and beyond*. Retrieved from <http://eprints.qut.edu.au/61773/>
- Castells, M. (2000). *The Rise of the Network Society*. New York: John Wiley & Sons.
- Chen, A. (2012, February 16). Inside Facebook's Outsourced Anti-Porn and Gore Brigade, Where 'Camel Toes' are More Offensive Than 'Crushed Heads'. Retrieved March 17, 2013, from <http://gawker.com/5885714/>
- Chen, G., & Rahman, F. (2008). Analyzing privacy designs of mobile social networking applications. *2008 IEEE/IFIP International Conference on Embedded and Ubiquitous Computing*. <http://doi.org/10.1109/EUC.2008.156>
- Chen, K. (2009, October 26). Memories of friends departed endure on Facebook. Retrieved April 27, 2015, from <https://www.facebook.com/notes/facebook/memories-of-friends-departed-endure-on-facebook/163091042130>
- Cheney-Lippold, J. (2011). A New Algorithmic Identity: Soft Biopolitics and the Modulation of Control. *Theory, Culture & Society*, 28(6), 164–181. <http://doi.org/10.1177/0263276411424420>
- Chun, W. H. K. (2006). *Control and freedom: Power and paranoia in the age of fiber optics*. Cambridge: MIT Press.
- Cirucci, A. M. (2015). Redefining privacy and anonymity through social networking affordances. *First Monday*, 20(7).
- Clegg, S., & Rhodes, C. (2006). Introduction: Questioning the ethics of management practice. In S. Clegg & C. Rhodes (Eds.), *Management ethics: Contemporary*

- contexts*. London: Routledge.
- Cole, M., & Avison, D. (2007). The potential of hermeneutics in information systems research. *European Journal of Information Systems*, 16(6), 820–833.
<http://doi.org/10.1057/palgrave.ejis.3000725>
- Collins, H. (1999). The TEA Set: Tacit knowledge and scientific networks. In M. Biagioli (Ed.), *The science studies reader*. New York: Psychology Press.
- Collins, P. H. (2002). *Black Feminist Thought*. London: Routledge.
- Colter, E. G., Hoffman, W., Pendleton, E., Redick, A., & Serlin, D. (Eds.). (1996). *Policing public sex*. Boston: South End Press.
- Consalvo, M. (2003). Zelda 64 and Video Game Fans: A Walkthrough of Games, Intertextuality, and Narrative. *Television & New Media*, 4(3), 321–334.
<http://doi.org/10.1177/1527476403253993>
- Cooper, D. (2012, October 12). “No fats or fems.” Retrieved December 3, 2014, from http://www.huffingtonpost.com/dale-cooper/grindr-discrimination_b_1948766.html
- Cox, G. (2013). *Speaking Code: Coding as Aesthetic and Political Expression*. Cambridge: MIT Press.
- Crary, J. (1988). Techniques of the Observer. *October*, 45, 3–35.
- Crawford, K., & Lumby, C. (2013). Networks of Governance: Users, Platforms, and the Challenges of Networked Media Regulation. *International Journal of Technology Policy and Law*, 2(1).
- Crook, J. (2013, October 2). Gay Gets Better (And More Targeted): Say Hello To The Next Generation Of Grindr. Retrieved December 1, 2013, from <http://techcrunch.com/2013/10/02/gay-gets-better-and-more-targeted-say-hello-to->

the-next-generation-of-grindr/

Crooks, R. N. (2013). The Rainbow Flag and the Green Carnation: Grindr in The Gay Village. *First Monday*, 18(11).

Davidson, J., & Martellozzo, E. (2013). Exploring Young People's Use of Social Networking Sites and Digital Media in the Internet Safety Context. *Information, Communication & Society*, 16(9), 1456–1476.

<http://doi.org/10.1080/1369118X.2012.701655>

Davis, J. L., & Jurgenson, N. (2014). Context collapse: theorizing context collusions and collisions. *Information, Communication & Society*, 17(4), 476–485.

<http://doi.org/10.1080/1369118X.2014.888458>

de Certeau, M. (1988). *The practice of everyday life*. Minnesota: University of Minnesota Press.

de Souza e Silva, A. (2006). From Cyber to Hybrid: Mobile Technologies as Interfaces of Hybrid Spaces. *Space and Culture*, 9(3), 261–278.

<http://doi.org/10.1177/1206331206289022>

de Souza e Silva, A., & Frith, J. (2010a). Locational Privacy in Public Spaces: Media Discourses on Location-Aware Mobile Technologies. *Communication, Culture & Critique*, 3(4), 503–525. <http://doi.org/10.1111/j.1753-9137.2010.01083.x>

de Souza e Silva, A., & Frith, J. (2010b). Locative Mobile Social Networks: Mapping Communication and Location in Urban Spaces. *Mobilities*, 5(4), 485–505.

<http://doi.org/10.1080/17450101.2010.510332>

de Souza e Silva, A., & Frith, J. (2012). *Mobile interfaces in public spaces: Locational privacy, control, and urban sociability*. New York: Routledge.

- Delany, S. (2001). *Times Square red, Times Square blue*. New York: New York University Press.
- Deleuze, G. (1992). Postscript on the Societies of Control. *October*, 59.
- Derian, Der, J. (2001). *Virtuous war: Mapping the military-industrial-media-entertainment network*. Boulder: Westview Press.
- Dessent, M. (2002). Browse-Wraps, Click-Wraps and Cyberlaw: Our Shrinking (Wrap) World. *Thomas Jefferson Law Review*, 25.
- Dibbell, J. (1991). *My tiny life: Crime and passion in a virtual world*. New York: Henry Holt.
- Dibbell, J. (1994). Rape in Cyberspace or How an Evil Clown, a Haitian Trickster Spirit, Two Wizards, and a Cast of Dozens Turned a Database into a Society. *Annual Survey of American Law*, 471–489.
- Donahue, B., & Stoner, J. (1997). The natural bears classification system: A classification system for bears and bearlike men version 1.10. In L. Wright (Ed.), *The bear book: Readings in the history and evolution of a gay male subculture*. New York: Harrington Park Press.
- Döring, N. M. (2009). The Internet's impact on sexuality: A critical review of 15 years of research. *Computers in Human Behavior*, 25(5), 1089–1101.
<http://doi.org/10.1016/j.chb.2009.04.003>
- Draper, N. (2014, July 14). *Reputation Inc.: The industrialization of digital self-presentation and online privacy*.
- Duguay, S. (2014). “He has a way gayer Facebook than I do”: Investigating sexual identity disclosure and context collapse on a social networking site. *New Media and*

Society. <http://doi.org/10.1177/1461444814549930>

Duncan, S. (2008). MySpace is also their space: Ideas for keeping children safe from sexual predators on social-networking sites. *Kentucky Law Journal*, 96.

Duyves, M. (1993). The minitel: the glittering future of a new invention. *Journal of Homosexuality*, 25(1-2), 193–203. http://doi.org/10.1300/J082v25n01_13

Easton, R. (2009, December 30). Will prudish app guidelines sink the iPhone? Retrieved March 17, 2013, from http://www.xtra.ca/public/National/Will_prudish_app_guidelines_sink_the_iPhone-7988.aspx

Ellison, N. B., Hancock, J. T., & Toma, C. L. (2012). Profile as promise: A framework for conceptualizing veracity in online dating self-presentations. *New Media and Society*, 14(1), 45–62. <http://doi.org/10.1177/1461444811410395>

Ellison, N., Heino, R., & Gibbs, J. (2006). Managing Impressions Online: Self-Presentation Processes in the Online Dating Environment. *Journal of Computer-Mediated Communication*, 11(2), 415–441. <http://doi.org/10.1111/j.1083-6101.2006.00020.x>

Elovitz, M. E., & Edwards, P. J. (1996). The D.O.H. papers: Regulating public sex in New York City. In E. G. Colter, W. Hoffman, E. Pendleton, A. Redick, & D. Serlin (Eds.), *Policing public sex*. Boston: South End Press.

Erlichmann, J. (2012, June 15). Grindr CEO on Gay Social App's Growth. Retrieved December 4, 2013, from <http://www.businessweek.com/videos/2012-06-15/grindr-ceo-on-gay-social-apps-growth>

Facebook. (n.d.). Facebook Community Standards. Retrieved March 17, 2013, from

<https://www.facebook.com/communitystandards>

Fanon, F. (1967). *Black Skin, White Masks*. London: Pluto Press.

Farman, J. (2012). *Mobile interface theory: Embodied space and locative media*. New York: Routledge.

Federal Trade Commission. (2009). *Negative Options*. Retrieved March 16, 2013, from <http://www.ftc.gov/os/2009/02/P064202negativeoptionreport.pdf>

Flanagan, M., Howe, D., & Nissenbaum, H. (2008). *Embodying values in technology: Theory and practice*. In J. van den Hoven & J. Weckert (Eds.), *Information technology and moral philosophy*. Cambridge: Cambridge University Press.

Flock, E. (2012, June 9). *Gay sex app Grindr goes political*. Retrieved April 6, 2014, from <http://www.usnews.com/news/blogs/washington-whispers/2012/09/06/gay-sex-app-grindr-goes-political>

Foucault, M. (1978a). *Discipline and Punish: The birth of the prison*. (A. Sheridan, Ed.). New York: Pantheon.

Foucault, M. (1978b). *The History of Sexuality: An Introduction*. New York: Pantheon.

Foucault, M. (1986). *The History of Sexuality: The care of the self*. (R. Hurley, Ed.). New York: Pantheon.

Friedman, B. (1996). *Value-sensitive design*. *Interactions*, 3(6).

Friedman, B., & Nathan, L. P. (2010). *Multi-lifespan information system design: a research initiative for the HCI community*. *CHI '10 Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.

Friedman, B., & Nissenbaum, H. (1996). *Bias in computer systems*. *ACM Transactions on Information Systems*, 14(3).

- Friedman, B., Nathan, L. P., Lake, M., & Grey, N. C. (2010). Multi-lifespan information system design in post-conflict societies: An evolving project in Rwanda. *CHI '10 Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.
- Frith, J. (2013). Turning life into a game: Foursquare, gamification, and personal mobility. *Mobile Media & Communication*, 1(2), 248–262.
<http://doi.org/10.1177/2050157912474811>
- Fuchs, C. (2011). An Alternative View of Privacy on Facebook. *Information*, 2(4), 140–165. <http://doi.org/10.3390/info2010140>
- Fuller, M. (2003). *Behind the blip: Essays on the culture of software*. New York: Autonomedia.
- Fuller, M. (2008). *Software Studies: A Lexicon*. Cambridge: MIT Press.
- Galison, P. (1997). *Image and Logic: A Material Culture of Microphysics*. Chicago: University of Chicago Press.
- Galloway, A. R. (2004). *Protocol*. Cambridge: The MIT Press.
- Gamson, J. (2009). *Freaks Talk Back*. Chicago: University of Chicago Press.
- Gandy, O. H. (1993). *The panoptic sort: A political economy of personal information*. Boulder: Westview Press.
- Gates, K. (2011). *Our Biometric Future*. New York: NYU Press.
- Gates, K. (2013). The cultural labor of surveillance: video forensics, computational objectivity, and the production of visual evidence. *Social Semiotics*, 23(2), 242–260.
<http://doi.org/10.1080/10350330.2013.777593>
- Gibson, J. J. (1986). *The theory of affordances: The ecological approach to visual perception*. London: L. Erlbaum.

- Giddens, A. (1990). *The Consequences of Modernity*. Stanford: Stanford University Press.
- Gillespie, T. (2007). *Wired Shut: Copyright and the Shape of Digital Culture*. Cambridge: MIT Press.
- Gillespie, T. (2010). The politics of “platforms.” *New Media and Society*, 12(3), 347–364. <http://doi.org/10.1177/1461444809342738>
- Gillespie, T. (2012, February 22). The dirty job of keeping Facebook clean. Retrieved January 30, 2013, from <http://socialmediacollective.org/2012/02/22/the-dirty-job-of-keeping-facebook-clean/>
- Gillette, F. (2013, February 7). Snapchat and the Erasable Future of Social Media. Retrieved November 26, 2014, from <http://www.businessweek.com/articles/2013-02-07/snapchat-and-the-erasable-future-of-social-media>
- Gilman, S. (1982). *Seeing the insane*. New York: John Wiley.
- Gilman, S. (1989). *Sexuality: An illustrated history*. New York: John Wiley.
- Gilman, S. (1995). *Health and illness: Images of difference*. London: Reaktion Books.
- Gindin, S. E. (2009). Nobody Reads Your Privacy Policy or Online Contract: Lessons Learned and Questions Raised by the FTC's Action against Sears. *Northwestern Journal of Technology and Intellectual Property*, 8.
- Glaser, B. G., & Strauss, A. L. (1967). *The Discovery of Grounded Theory*. Mill Valley: Sociology Press.
- Goffman, E. (1959). *The Presentation of Self in Everyday Life*. New York: Anchor Doubleday.
- Goldman, E. (2015, March 19). Online Dating App Grindr Isn't Liable For Underage

- 'Threesome'. Retrieved April 20, 2015, from <http://www.forbes.com/sites/ericgoldman/2015/03/19/online-dating-app-grindr-isnt-liable-for-underage-threesome/>
- Gomez, J., Pinnick, T., & Soltani, A. (2009). KnowPrivacy. *University of California Berkeley School of Information*.
- Google. (n.d.). Google Play developer program policies. Retrieved December 7, 2012, from <http://play.google.com/about/developer-content-policy.html>
- Gosine, A. (2007). Brown to blonde at Gay.com: Passing white in queer cyberspace. In K. O'Riordan & D. J. Philips (Eds.), *Queer online: Media, technology, and sexuality*. New York: Peter Lang Pub Inc.
- Graham, S. (2004). Vertical geopolitics: Baghdad and after. *Antipode*, 36(1), 12–23.
- Gremore, G. (2014, October 23). Man Assaulted With An Electric Cattle Prod In Grindr Hookup Gone Horribly Wrong. Retrieved April 20, 2015, from <http://www.queerty.com/man-assaulted-with-an-electric-cattle-prod-in-grindr-hookup-gone-horribly-wrong-20141023>
- Grindr. (2012, May 9). Grindr for Equality: Election 2012. Retrieved April 6, 2014, from <http://grindr.com/blog/grindr-for-equality-election-2012>
- Grindr. (2013a, July 12). Grindr Privacy Policy. Retrieved December 2, 2013, from <http://grindr.com/privacy-policy>
- Grindr. (2013b, September 26). Grindr fact sheet. Retrieved March 25, 2014, from <http://grindr.com/press>
- Grindr. (2014a, June 4). Grindr terms of service. Retrieved December 3, 2014, from <http://grindr.com/terms-of-service>

- Grindr. (2014b, August 28). Grindr security. Retrieved December 2, 2014, from <http://grindr.com/blog/grindr-security>
- Grindr. (2014c, September 5). Grindr's location security update. Retrieved December 2, 2014, from <http://grindr.com/blog/grindr-location-security-update>
- Grindr. (2014d, November 12). Grindr safety tips. Retrieved December 3, 2014, from <https://help.grindr.com/hc/en-us/articles/200302014-Grindr-Safety-Tips>
- Grindr. (n.d.). Grindr profile guidelines. Retrieved December 7, 2012, from <http://grindr.com/profile-guidelines>
- Grinnell, C. K. (2009). From Consumer to Prosumer to Producer: Who Keeps Shifting My Paradigm? (We Do!). *Public Culture*, 21(3), 577–598.
<http://doi.org/10.1215/08992363-2009-009>
- Gross, L. (1993). *Contested closets: The politics and ethics of outing*. Minneapolis: University of Minnesota Press.
- Gross, L. (2013). *Up from Invisibility*. New York: Columbia University Press.
- Gudelunas, D. (2012a). Generational Differences Among Gay Men and Lesbians: Social and Media Change. Presented at the International Communication Association Virtual Preconference.
- Gudelunas, D. (2012b). There's an App for that: The Uses and Gratifications of Online Social Networks for Gay Men. *Sexuality & Culture*, 16(4).
<http://doi.org/10.1007/s12119-012-9127-4>
- Habib, C. (2013, June 26). Facing the torsos. *The Stranger*. Retrieved from <http://www.thestranger.com/seattle/Content?oid=17107224>
- Hacking, I. (2006). Making up people. *London Review of Books*, 28(18), 23–26.

- Hafertepen, D. (n.d.). Gay Cliques Census. Retrieved November 25, 2012, from <http://studiomoh.com/fun/census/>
- Halavais, A. (2009). *Search Engine Society*. Cambridge: Polity.
- Halberstam, J. (1998). *Female masculinity*. Durham: Duke University Press.
- Hall, L. (2015, March 3). Sydney man jailed after luring then robbing victims he met on gay social networking app Grindr. Retrieved April 20, 2015, from <http://www.smh.com.au/nsw/sydney-man-jailed-after-luring-then-robbing-victims-he-met-on-gay-social-networking-app-grindr-20150303-13tqny.html>
- Hall, S. (1973). Encoding/Decoding. In *Culture, Media, Language: Working papers in cultural studies, 1972-79*. Oxford: Psychology Press.
- Han, C.-S. (2007). They Don't Want To Cruise Your Type: Gay Men of Color and the Racial Politics of Exclusion. *Social Identities, 13*(1), 51–67. <http://doi.org/10.1080/13504630601163379>
- Hancock, J. T., & Toma, C. L. (2009). Putting Your Best Face Forward: The Accuracy of Online Dating Photographs. *Journal of Computer-Mediated Communication, 59*(2), 367–386. <http://doi.org/10.1111/j.1460-2466.2009.01420.x>
- Hands, J. (2013). Introduction: Politics, Power and “Platformivity.” *Culture Machine, 14*(1), 1–9.
- Haraway, D. J. (1991). *Simians, Cyborgs, and Women*. New York: Routledge.
- Harding, S. G. (1991). *Whose Science? Whose Knowledge?* Ithaca: Cornell University Press.
- Hayles, N. K. (1999). *How we became posthuman: Virtual bodies in cybernetics, literature, and informatics*. Chicago: University of Chicago Press.

- Hayles, N. K. (2004). Print Is Flat, Code Is Deep: The Importance of Media-Specific Analysis. *Poetics Today*, 25(1), 67–90. <http://doi.org/10.1215/03335372-25-1-67>
- Hebdige, D. (1979). *Subculture: The meaning of style*. London: Routledge.
- Hestres, L. (2013). App Neutrality: Apple's App Store and Freedom of Expression Online. *International Journal of Communication*, 7.
- Hine, C. (2000). *Virtual Ethnography*. London: SAGE.
- Hongladarom, S. (2011). Personal Identity and the Self in the Online and Offline World. *Minds and Machines*, 21(4), 533–548. <http://doi.org/10.1007/s11023-011-9255-x>
- Humphreys, Laud. (1975). *Tearoom Trade*. Chicago: Aldine De Gruyter.
- Humphreys, Lee. (2007). Mobile Social Networks and Social Practice: A Case Study of Dodgeball. *Journal of Computer-Mediated Communication*, 13(1), 341–360. <http://doi.org/10.1111/j.1083-6101.2007.00399.x>
- Humphreys, Lee. (2012). Connecting, Coordinating, Cataloguing: Communicative Practices on Mobile Social Networks. *Journal of Broadcasting & Electronic Media*, 56(4), 494–510. <http://doi.org/10.1080/08838151.2012.732144>
- Hutchby, I. (2001). Technologies, texts and affordances. *Sociology*, 35(2), 441–456.
- Ibrahim, Y. (2010). The Breastfeeding Controversy and Facebook. *International Journal of E-Politics*, 1(2). <http://doi.org/10.4018/jep.2010040102>
- Ingebretsen, E. (1999). Gone shopping: The commercialization of same-sex desire. *International Journal of Sexuality and Gender Studies*, 4(2), 125–148.
- Israeli, H. B. (1995). From ‹Bonehead› to ‹cLoNehEAd› : Nicknames, play and identity on internet relay chat. *Journal of Computer-Mediated Communication*, 1(2).

<http://doi.org/10.1111/jcmc.1995.1.issue-2/issuetoc>

Jackson, J. L., Jr. (2005). *Real Black: Adventures in racial sincerity*. Chicago: University of Chicago Press.

Jenkins, H. (2006). *Convergence culture*. New York: NYU Press.

Joinson, A. N. (2008). Looking at, looking up or keeping up with people?: motives and use of facebook. *Chi '08*, 1027–1036. <http://doi.org/10.1145/1357054.1357213>

Jurgenson, N. (2012). When Atoms Meet Bits: Social Media, the Mobile Web and Augmented Revolution. *Future Internet*, 4(4), 83–91.

<http://doi.org/10.3390/fi4010083>

Kaplan, C. (2006). Precision Targets: GPS and the Militarization of U.S. Consumer Identity. *American Quarterly*, 58(3), 693–714. <http://doi.org/10.1353/aq.2006.0061>

Kapp, M. (2011, May 27). Grindr: Welcome to the World's Biggest, Scariest Gay Bar.

Vanity Fair. Retrieved from

<http://www.vanityfair.com/culture/features/2011/05/grindr-201105.print>

Kendall, L. (1998). Meaning and identity in “cyberspace”: The performance of gender, class, and race online. *Symbolic Interaction*, 21(2).

<http://doi.org/10.1525/si.1998.21.2.129/full>

Kennedy, H. (2011). *Net Work*. London: Palgrave Macmillan.

Kleeman, S. (2015, October 1). Instagram Finally Revealed the Reason It Banned

Nipples — It's Apple. Retrieved October 4, 2015, from

<http://mic.com/articles/126137/instagram-banned-nipples-because-of-apple>

Koskela, H. (2002). “Cam Era”—the contemporary urban Panopticon. *Surveillance & Society*, 1(3), 292–313.

- Langlois, G. (2013). Participatory Culture and the New Governance of Communication: The Paradox of Participatory Media. *Television & New Media*, 14(2), 91–105.
<http://doi.org/10.1177/1527476411433519>
- Langlois, G., McKelvey, F., Elmer, G., & Werbin, K. (2009). Mapping commercial Web 2.0 worlds: Towards a new critical ontogenesis. *Fibreculture*, 14.
- LaRosa, P., & Cramer, M. (2009). Seven days of rage: The deadly crime spree of the Craigslist Killer. New York: Simon & Schuster.
- Latour, B. (2005). Reassembling the Social. Oxford: Oxford University Press.
- Lattanzio, V. (2014, April 11). Tourist Raped, Beaten & Robbed After Meeting Man on Gay Dating App Grindr. Retrieved April 20, 2015, from
<http://www.nbcphiladelphia.com/news/local/Tourist-Sexually-Assaulted-Beaten--Robbed-After-Meeting-Suspect-on-GRINDR-App-254841611.html>
- Lebduska, L. (2014). Emoji, Emoji, What for Art Thou? *Harlot: a Revealing Look at the Arts of Persuasion*.
- Lehmiller, J. J., & Ioerger, M. (2014). Social Networking Smartphone Applications and Sexual Health Outcomes among Men Who Have Sex with Men. *PLoS ONE*, 9(1).
<http://doi.org/10.1371/journal.pone.0086603>
- Lessig, L. (1999). Code: And other laws of cyberspace. New York: Basic Books.
- Levine, M. (1998). Gay macho: The life and death of the homosexual clone. New York: New York University Press.
- Licoppe, C. (2004). “Connected” presence: the emergence of a new repertoire for managing social relationships in a changing communication technoscape. *Environment and Planning D: Society and Space*, 22(1), 135–156.

<http://doi.org/10.1068/d323t>

Light, B. (2014). *Disconnecting with Social Networking Sites*. London: Palgrave Macmillan.

Light, B., Fletcher, G., & Adam, A. (2008). Gay men, Gaydar and the commodification of difference. *Information Technology & People*, 21(3), 300–314.

<http://doi.org/10.1108/09593840810896046>

Livia, A. (2002). Public and clandestine: gay men's pseudonyms on the French Minitel. *Sexualities*, 5(2), 201–217.

Lustig, N. (2012, June 6). 2.89m Facebook users will die in 2012, 580,000 in the USA.

Retrieved April 27, 2015, from <http://www.nathanlustig.com/tag/how-many-facebook-users-die-per-year/>

Lyon, D. (2001). *Surveillance society: Monitoring everyday life*. Buckingham: Open University Press.

MacKenzie, A. (2006). *Cutting Code: Software And Sociality*. New York: Peter Lang.

Madejski, M., Johnson, M. L., & Bellovin, S. M. (2011). The failure of online social network privacy settings. Retrieved November 5, 2013, from

<http://academiccommons.columbia.edu/catalog/ac:135406>

Madison, M. J. (1998). Legal-Ware: Contract and Copyright in the Digital Age. *Fordham Law Review*, 67, 1025–1143.

Majchrzak, A., Faraj, S., Kane, G. C., & Azad, B. (2013). The Contradictory Influence of Social Media Affordances on Online Communal Knowledge Sharing. *Journal of Computer-Mediated Communication*, 19(1), 38–55.

<http://doi.org/10.1111/jcc4.12030>

- Manhunt. (2009a, January 14). Picture guidelines. Retrieved December 7, 2012, from <http://help.manhunt.net/question.php?ID=79>
- Manhunt. (2009b, January 14). What's not allowed. Retrieved December 7, 2012, from <http://help.manhunt.net/question.php?ID=81>
- Manhunt. (2009c, March 23). Terms of access and use. Retrieved December 7, 2012, from <http://help.manhunt.net/question.php?ID=225>
- Manhunt. (2011, March 17). App approved picture guidelines. Retrieved December 7, 2012, from <http://help.manhunt.net/question.php?ID=339>
- Mann, S. (2005). Sousveillance and cyborglogs: a 30-year empirical voyage through ethical, legal, and policy issues. *Presence, 14*(6), 625–646.
- Mann, S., Nolan, J., & Wellman, B. (2003). Sousveillance: Inventing and Using Wearable Computing Devices for Data Collection in Surveillance Environments. *Surveillance & Society, 1*(3), 331–355.
- Manovich, L. (2001). *The Language of New Media*. Cambridge: MIT Press.
- Manovich, L. (2006). The poetics of augmented space. *Visual Communication, 5*(2), 219–240. <http://doi.org/10.1177/1470357206065527>
- Manovich, L. (2013). *Software takes command: Extending the language of new media*. New York: Bloomsbury.
- Marcus, G. E. (1995). Ethnography in/of the world system: the emergence of multi-sited ethnography. *Annual Review of Anthropology, 24*. <http://doi.org/10.2307/2155931>
- Marwick, A. (2012). The public domain: Social surveillance in everyday life. *Surveillance & Society, 9*(4), 378–393.
- Marwick, A. E. (2005). Selling your self: Online identity in the age of a commodified

internet.

Marwick, A. E. (2013). *Status Update*. New Haven: Yale University Press.

Marwick, A. E., & boyd, D. (2011). I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media and Society*, *13*(1), 114–133. <http://doi.org/10.1177/1461444810365313>

Massad, J. (2002). Re-orienting desire: The gay international and the Arab world. *Public Culture*, *14*(2), 361–385.

Mateljan, F. (2014, December 1). City attorney Feuer secures conviction under state's 'revenge porn' law. Retrieved from http://atty.lacity.org/stellent/groups/electedofficials/@atty_contributor/documents/contributor_web_content/lacityp_029467.pdf

May, L. (2015, January 31). Grindr in the Kremlin: Gay and Online in Putin's Russia. Retrieved April 20, 2015, from <http://globalvoicesonline.org/2015/01/31/grindr-in-the-kremlin-gay-and-online-in-putins-russia/>

Mayer-Schönberger, V. (2009). *Delete*. Princeton: Princeton University Press.

Maynard, S. (1994). Through a hole in the lavatory wall: Homosexual subcultures, police surveillance, and the dialectics of discovery, Toronto, 1890-1930. *Journal of the History of Sexuality*, *5*(2), 207–242.

McVeigh-Schultz, J., & Baym, N. K. (2015). Thinking of You: Vernacular Affordance in the Context of the Microsocial Relationship App, Couple. *Social Media + Society*, *1*(2), 1–13. <http://doi.org/10.1177/2056305115604649>

Meyrowitz, J. (1985). *No Sense of Place: The Impact of Electronic Media on Social Behavior*. Oxford: Oxford University Press.

- Miller, D., & Slater, D. (2000). *The Internet: an ethnographic approach*. Berg Publishers.
- Miller, M. E. (2015, April 29). N.D. legislator who voted against gay rights bill caught sending pics on Grindr. Retrieved April 30, 2015, from <http://www.washingtonpost.com/news/morning-mix/wp/2015/04/29/n-d-legislator-caught-sending-grindr-pics-after-voting-against-gay-rights-bill/>
- Miller, T. (1993). *The well-tempered self*. Baltimore: Johns Hopkins University Press.
- Minassian, A. (1997). Death of Copyright: Enforceability of Shrinkwrap Licensing Agreements, *The. UCLA Law Review*, 45, 569–609.
- Molinet, J. (2014, September 15). HIV-positive teacher uses Grindr for teen hookup: cops. Retrieved April 20, 2015, from <http://www.nydailynews.com/news/crime/hiv-positive-teacher-grindr-teen-hookup-cops-article-1.1939646>
- Mosendz, P. (2015, January 30). Gay Dating App Scruff Puts 48-Foot Billboard Near Super Bowl Stadium, Sees Usage Uptick. Retrieved February 10, 2015, from <http://www.newsweek.com/gay-dating-app-scruff-puts-48-foot-billboard-near-super-bowl-stadium-sees-303355>
- Mowlabocus, S. (2010a). *Gaydar culture: Gay men, technology and embodiment in the digital age*. Surrey: Ashgate.
- Mowlabocus, S. (2010b). Look at Me!: Images, validation, and cultural currency on Gaydar. In C. Pullen & M. Cooper (Eds.), *LGBT identity and online new media*. London: Routledge.
- Mowlabocus, S. (2014a, September 3). Grindr relents to backlash – but does it really respect its users? Retrieved December 2, 2014, from <http://theconversation.com/grindr-relents-to-backlash-but-does-it-really-respect-its->

users-31198

- Mowlabocus, S. (2014b, November 8). Grindr's locator "glitch" was a major fail. It revealed the company's lack of empathy for its gay users. Retrieved April 20, 2015, from <http://www.washingtonpost.com/posteverything/wp/2014/09/08/grindr-locator-glitch-was-a-major-fail-it-revealed-the-companys-lack-of-empathy-for-its-gay-users/>
- Moylan, B. (2011, July 20). The founder of Grindr is just as addicted as everyone else. Retrieved December 3, 2014, from <http://gawker.com/5823107/the-founder-of-grindr-is-just-as-addicted-as-everyone-else>
- Nafus, D., & Sherman, J. (2014). This One Does Not Go Up To 11: The Quantified Self Movement as an Alternative Big Data Practice. *International Journal of Communication*, 8, 1784–1794.
- Nagy, P., & Neff, G. (2015). Imagined Affordance: Reconstructing a Keyword for Communication Theory. *Social Media + Society*, 1(2), 1–9.
<http://doi.org/10.1177/2056305115603385>
- Nakamura, L. (2001). Race in/for cyberspace: Identity tourism and racial passing on the internet. In D. Trend (Ed.), *Reading digital culture*. Malden: Blackwell.
- Nakamura, L. (2002). *Cybertypes: Race, ethnicity, and identity on the internet*. New York: Routledge.
- Nathan, L. P., Lake, M., Grey, N. C., & Nilsen, T. (2011). Multi-lifespan information system design: Investigating a new design approach in Rwanda. *Science, Technology, & Human Values*, 591–597.
- Negroponte, N. (1995). *Being Digital*. New York: Vintage.

- Nelson, D. M. (2001). Stumped identities: Body image, bodies politic, and the Mujer Maya as prosthetic. *Cultural Anthropology*, 16(3), 314–353.
- Nissenbaum, H. (2001). How computer systems embody values. *Computer*, 34(3).
- Nissenbaum, H. F. (2010). *Privacy in context*. Stanford: Stanford University Press.
- North, A. (2014, November 19). How OkCupid Has Become More Inclusive on Gender and Sexuality. Retrieved September 27, 2015, from http://op-talk.blogs.nytimes.com/2014/11/19/how-okcupid-has-become-more-inclusive-on-gender-and-sexuality/?_r=0
- O'Bryan, W. (2012, August 16). Zero feet away. Retrieved August 28, 2012, from <http://www.metroweekly.com/feature/?ak=7657&pagenumber=all>
- O'Riordan, K. (2005). From usenet to Gaydar: a comment on queer online community. *ACM SigGroup Bulletin*, 25(2), 28–32.
- O'Riordan, K. (2007). Queer theories and cybersubjects: Intersecting figures. In K. O'Riordan & D. J. Philips (Eds.), *Queer online: Media, technology, and sexuality*. New York: Peter Lang Pub Inc.
- O'Riordan, K., & Philips, D. J. (2007). *Queer online: Media, technology, and sexuality*. New York: Peter Lang.
- Obar, J. (2013). Phantom data sovereigns: Walter Lippmann, big data and the fallacy of personal data sovereignty. Retrieved November 12, 2014, from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2239188
- Ortner, S. B. (2010). Access: Reflections on studying up in Hollywood. *Ethnography*, 11(2), 211–233. <http://doi.org/10.1177/1466138110362006>
- Oudshoorn, N., & Pinch, T. J. (2003). *How Users Matter*. Cambridge: MIT Press.

- Papacharissi, Z. (2002a). The presentation of self in virtual life: Characteristics of personal home pages. *Journalism & Mass Communication Quarterly*, 79(3), 643–660.
- Papacharissi, Z. (2002b). The Self Online: The Utility of Personal Home Pages. *Journal of Broadcasting & Electronic Media*, 46(3), 346–368.
http://doi.org/10.1207/s15506878jobem4603_3
- Papacharissi, Z. (2009). The virtual geographies of social networks: a comparative analysis of Facebook, LinkedIn and ASmallWorld. *New Media and Society*, 11(1-2), 199–220. <http://doi.org/10.1177/1461444808099577>
- Papacharissi, Z. (2012). Without you, I'm nothing: Performances of the self on Twitter. *International Journal of Communication*, 6, 1989–2006.
- Parks, L. (2001). Plotting the personal: Global Positioning Satellites and interactive media. *Cultural Geographies*, 8(2), 209–222.
<http://doi.org/10.1177/096746080100800205>
- Patton, C. (1990). *Inventing AIDS*. New York: Routledge.
- Paul, J. P., Ayala, G., & Choi, K.-H. (2010). Internet Sex Ads for MSM and Partner Selection Criteria: The Potency of Race/Ethnicity Online. *The Journal of Sex Research*, 47(6), 528–538. <http://doi.org/10.1080/00224490903244575>
- Pike, G. (2004). Shrink-wrap, Click-wrap, Now Browse-wrap. *Information Today*, 21(3).
- Poster, M. (1995). *The second media age*. Cambridge: Polity Press.
- Pulkkinen, L. (2014, May 12). Police: Grindr sex app date ended in hammer attack, biting. Retrieved April 20, 2015, from <http://www.seattlepi.com/local/article/Police-Grindr-sex-app-date-ended-in-hammer-5473093.php>

- Queerty. (2010, April 14). 16 Grindr Profiles Now Banned Under the New Puritan Rules. Retrieved March 17, 2013, from <http://www.queerty.com/16-grindr-profiles-now-banned-under-the-new-puritan-rules-20100414/>
- Race, K. (2015). "Party and Play": Online hook-up devices and the emergence of PNP practices among gay men. *Sexualities*, 18(3), 253–275. <http://doi.org/10.1177/1363460714550913>
- Raj, S. (2011). Grinding bodies: Racial and affective economies of online queer desire. *Critical Race and Whiteness Studies*, 7(2).
- ReadWrite Editors. (2014, February 13). Facebook Provides 56 New Gender Identity Options. Retrieved March 31, 2015, from <http://readwrite.com/2014/02/13/facebook-provides-50-new-gender-identity-options>
- Rendina, H. J., Jimenez, R. H., Grov, C., Ventuneac, A., & Parsons, J. T. (2013). Patterns of Lifetime and Recent HIV Testing Among Men Who Have Sex with Men in New York City Who Use Grindr. *AIDS and Behavior*, 18(1), 41–49. <http://doi.org/10.1007/s10461-013-0573-2>
- Rheingold, H. (2000). *The virtual community: Homesteading on the electronic frontier*. Cambridge: MIT Press.
- Rice, E., Holloway, I., & Winetrobe, H. (2012). Sex risk among young men who have sex with men who use Grindr, a smartphone geosocial networking application. *Journal of AIDS & Clinical Research*. <http://doi.org/10.4172/2155-6113.S4-004>
- Riggs, D. W. (2012). Anti-Asian Sentiment Amongst a Sample of White Australian Men on Gaydar. *Sex Roles*, 68(11-12), 768–778. [http://doi.org/10.1007/s11199-012-0119-](http://doi.org/10.1007/s11199-012-0119-5)

- Ritzer, G., & Jurgenson, N. (2010). Production, Consumption, Prosumption: The nature of capitalism in the age of the digital “prosumer.” *Journal of Consumer Culture*, 10(1), 13–36. <http://doi.org/10.1177/1469540509354673>
- Rogers, E. M. (2003). *Diffusion of Innovations*, 5th Edition. New York: Simon and Schuster.
- Rosser, B. R. S., Wilkerson, J. M., Smolenski, D. J., Oakes, J. M., Konstan, J., Horvath, K. J., et al. (2011). The Future of Internet-Based HIV Prevention: A Report on Key Findings from the Men’s INternet (MINTS-I, II) Sex Studies. *AIDS and Behavior*, 15(S1), 91–100. <http://doi.org/10.1007/s10461-011-9910-5>
- Roth, Y. (2014). Locating the “Scruff Guy”: Theorizing Body and Space in Gay Geosocial Media. *International Journal of Communication*, 8, 2113–2133.
- Schäfer, M. T. (2011). *Bastard Culture! How User Participation Transforms Cultural Production*. Amsterdam: Amsterdam University Press.
- Schofield, J. (2014, October 30). What happens to your Facebook account when you die? Retrieved April 27, 2015, from <http://www.theguardian.com/technology/askjack/2014/oct/30/what-happens-to-your-facebook-account-when-you-or-a-loved-one-dies>
- Schrock, A. R. (2015). Communicative Affordances of Mobile Media: Portability, Availability, Locatability, and Multimediality. *International Journal of Communication*, 9.
- Schwartz, P. M. (2004). Property, privacy, and personal data. *Harvard Law Review*, 117(7), 2056–2128.
- Schwartz, R., & Halegoua, G. R. (2014). The spatial self: Location-based identity

- performance on social media. *New Media and Society*.
<http://doi.org/10.1177/1461444814531364>
- Scott, J. C. (1999). *Seeing Like a State*. New Haven: Yale University Press.
- Scruff. (2012, November). Terms of service. Retrieved December 7, 2012, from
<http://www.scruffapp.com/en/tos/>
- Scruff. (n.d.). Profile guidelines. Retrieved December 7, 2012, from
<http://www.scruffapp.com/en/guidelines/>
- Seidman, S. (1992). *Embattled eros: Sexual politics and ethics in contemporary America*.
New York: Routledge.
- Sekula, A. (1986). The body and the archive. *October*, 39, 3–64.
- Sender, K. (2003). Sex Sells: Sex, class, and taste in commercial gay and lesbian media.
GLQ: a Journal of Lesbian and Gay Studies, 9(3), 331–365.
- Sender, K. (2004). *Business, not politics: The making of the gay market*. New York:
Columbia University Press.
- Sessions, L. F. (2009). “You Looked Better on MySpace”: Deception and authenticity on
the Web 2.0. *First Monday*, 14(7).
- Shafer, B. J., & Adams, A. E. (2005). Jurisprudence of Doubt: Obscenity, Indecency, and
Morality at the Dawn of the 21st Century. *Michigan Bar Journal*, 84(6).
- Shaw, D. F. (1997). Gay men and computer communication: A discourse of sex and
identity in cyberspace. In S. G. Jones & S. Jones (Eds.), *Virtual culture: Identity and
communication in cybersociety*. London: SAGE Publications.
- Shein, E. (2013). Ephemeral data. *Communications of the ACM*, 56(9), 20–22.
<http://doi.org/10.1145/2500468.2500474>

- Slater, D. (1998). Trading Sexpics on IRC: Embodiment and Authenticity on the Internet. *Body & Society*, 4(4), 91–117. <http://doi.org/10.1177/1357034X98004004005>
- Slater-Robins, M. (2015, September 30). Instagram's CEO admitted the reason it censors some photos of female nipples from the app is to keep Apple happy. Retrieved October 4, 2015, from <http://www.businessinsider.com/why-instagram-bans-freethenipple-2015-9?r=UK&IR=T>
- Social Media App May Have Played Part in Alleged Sexual Assault. (2015, April 7). Social Media App May Have Played Part in Alleged Sexual Assault. Retrieved April 20, 2015, from krv.com/news/local-news/Social-Media-App-May-Have-Played-Part-in-Alleged-Sexual-Assault/32243008#
- Solove, D. J. (2006). *The Digital Person: Technology and Privacy in the Information Age*. New York: NYU Press.
- Solove, D. J. (2008). *The Future of Reputation: Gossip, Rumor, and Privacy on the Internet*. New Haven: Yale University Press.
- Statt, N. (2015, October 16). The company behind Tinder and OkCupid just filed to go public. Retrieved October 17, 2015, from <http://www.theverge.com/2015/10/16/9558141/match-group-tinder-okcupid-ipo-iac>
- Stone, A. R. (1991). Will the real body please stand up? In M. Benedikt (Ed.). Cambridge: MIT Press.
- Stone, A. R. (1995). *The war of desire and technology at the close of the mechanical age*. Cambridge: MIT Press.
- Stroud, S. R. (2014). The Dark Side of the Online Self: A Pragmatist Critique of the Growing Plague of Revenge Porn. *Journal of Mass Media Ethics*, 29(3), 168–183.

<http://doi.org/10.1080/08900523.2014.917976>

Stutzman, F., & Kramer-Duffield, J. (2010). Friends only: examining a privacy-enhancing behavior in facebook. *CHI '10 Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.

Suckling, L. (2014, February 19). The lure of finding love by location. Retrieved December 3, 2014, from

http://www.nzherald.co.nz/lifestyle/news/article.cfm?c_id=6&objectid=11204189

Sundén, J. (2003). *Material Virtualities*. New York: Peter Lang.

Sutko, D. M., & de Souza e Silva, A. (2011). Location-aware mobile media and urban sociability. *New Media and Society*, *13*(5), 807–823.

<http://doi.org/10.1177/1461444810385202>

Swan, M. (2012). Sensor Mania! The Internet of Things, Wearable Computing, Objective Metrics, and the Quantified Self 2.0. *Journal of Sensor and Actuator Networks*, *1*(3), 217–253. <http://doi.org/10.3390/jsan1030217>

Tagg, J. (1988). *The burden of representation: Essays on photographs and histories*. Amherst: University of Massachusetts Press.

Tasker, T., & Pakcyk, D. (2008). Cyber-Surfing on the High Seas of Legalese: Law and Technology of Internet Agreements. *Albany Law Journal of Science and Technology*, *18*, 79–149.

Terranova, T. (2000). Free labor: Producing culture for the digital economy. *Social Text*, *18*(2).

Terranova, T. (2004). *Network culture*. London: Pluto Press.

Tharrett, M. (2014, October 4). Grindr Bans Glass Box Performance Artist After He's

- Attacked By Angry Visitor. Retrieved November 26, 2014, from <http://www.queerty.com/grindr-bans-glass-box-performance-artist-after-hes-attacked-by-angry-visitor-20141004>
- Thomas, J. (2011, June 30). The gay bar: Its new competition. Retrieved March 21, 2013, from http://www.slate.com/articles/life/the_gay_bar/2011/06/the_gay_bar_2.html
- Tokunaga, R. S. (2011). Social networking site or social surveillance site? Understanding the use of interpersonal electronic surveillance in romantic relationships. *Computers in Human Behavior*, 27(2), 705–713. <http://doi.org/10.1016/j.chb.2010.08.014>
- Toma, C. L., Hancock, J. T., & Ellison, N. B. (2008). Separating Fact From Fiction: An Examination of Deceptive Self-Presentation in Online Dating Profiles. *Personality and Social Psychology Bulletin*, 34(8), 1023–1036. <http://doi.org/10.1177/0146167208318067>
- Towle, A. (2010, April 14). Grindr's new terms of service (for puritans). Retrieved April 8, 2013, from <http://www.towleroad.com/2010/04/read-grindr-s-new-terms-of-service-for-puritans.html>
- Trebay, G. (2014, December 12). The Sex Education of Grindr's Joel Simkhai. Retrieved February 10, 2015, from http://www.nytimes.com/2014/12/14/fashion/the-sex-education-of-grindr-s-joel-simkhai.html?_r=0
- Trilling, L. (1971). *Sincerity and authenticity*. Cambridge: Harvard University Press.
- Tsjeng, Z. (2014, October 3). Artist causes outrage with public broadcast of Grindr PMs. Retrieved November 26, 2014, from <http://www.dazeddigital.com/artsandculture/article/22063/1/artist-causes-outrage-with-public-broadcast-of-grindr-pms>

- Turkle, S. (1994). Constructions and reconstructions of self in virtual reality: Playing in the MUDs. *Mind*, 1(3), 158–167. <http://doi.org/10.1080/10749039409524667>
- Turkle, S. (1995). *Life on the screen: Identity in the age of the internet*. New York: Simon and Schuster.
- Turkle, S. (2011). *Alone Together: Why We Expect More from Technology and Less from Each Other*. New York: Basic Books.
- Turner, F. (2005). Where the counterculture met the new economy: The WELL and the origins of virtual community. *Technology and Culture*, 46.
- Turner, F. (2006). *From Counterculture to Cyberculture: Stewart Brand, the Whole Earth Network, and the Rise of Digital Utopianism*. Chicago: University of Chicago Press.
- Turow, J. (2003). *Americans & Online Privacy: The System is Broken*. Annenberg Public Policy Center.
- Turow, J. (2012). *The Daily You: How the New Advertising Industry Is Defining Your Identity and Your Worth*. New Haven: Yale University Press.
- Turow, J., & Draper, N. (2012). Advertising's new surveillance ecosystem. In D. Lyon, K. Ball, & K. Haggerty (Eds.), *Routledge Handbook of Surveillance Studies*. London.
- Turow, J., King, J., Hoofnagle, C., Bleakley, A., & Hennessy, M. (2009). *Americans reject tailored advertising and three activities that enable it*. Available at SSRN 1478214. Annenberg Public Policy Center.
- Urry, J. (2002). Mobility and Proximity. *Sociology*, 36(2), 255–274. <http://doi.org/10.1177/0038038502036002002>
- Van De Wiele, C., & Tong, S. T. (2014). Breaking boundaries: the uses & gratifications of grindr (pp. 619–630). Presented at the the 2014 ACM International Joint

- Conference, ACM Press. <http://doi.org/10.1145/2632048.2636070>
- van Dijck, J. (2009). Users like you? Theorizing agency in user-generated content. *Media, Culture & Society*, 31(1). <http://doi.org/10.1177/0163443708098245>
- van Dijck, J. (2013a). Facebook and the engineering of connectivity: A multi-layered approach to social media platforms. *Convergence: the International Journal of Research Into New Media Technologies*, 19(2), 141–155. <http://doi.org/10.1177/1354856512457548>
- van Dijck, J. (2013b). *The culture of connectivity: A critical history of social media*. Oxford: Oxford University Press.
- van Dijck, J. (2013c). “You have one identity”: performing the self on Facebook and LinkedIn. *Media, Culture & Society*, 35(2), 199–215. <http://doi.org/10.1177/0163443712468605>
- van Dijck, J., & Nieborg, D. (2009). Wikinomics and its discontents: a critical analysis of Web 2.0 business manifestos. *New Media and Society*, 11(5), 855–874. <http://doi.org/10.1177/1461444809105356>
- van Dijck, J., & Poell, T. (2013). Understanding social media logic. *Media and Communication*, 1(1), 2–14. <http://doi.org/10.12924/mac2013.01010002>
- Vaz, P., & Bruno, F. (2003). Types of Self-Surveillance: from abnormality to individuals' at risk'. *Surveillance & Society*, 1(3), 272–291.
- Verhoeven, D. (2014, October 10). “Wanna Play?” A reflection by Dries Verhoeven. Retrieved November 26, 2014, from [http://www.driesverhoeven.com/sites/default/files/uploads/wanna_play-
een_reflectie_eng_def.pdf](http://www.driesverhoeven.com/sites/default/files/uploads/wanna_play-
een_reflectie_eng_def.pdf)

- Vernon, P. (2010, July 3). Grindr: a new sexual revolution? *The Observer*. Retrieved from <http://www.guardian.co.uk/media/2010/jul/04/grindr-the-new-sexual-revolution/print>
- Villarreal, D. (2011, July 19). Now There's A Place To Publicly Humiliate Grindr Douchebags. Retrieved December 3, 2014, from <http://www.queerty.com/now-theres-a-place-to-publicly-humiliate-grindr-douchebags-20110719>
- Wakeford, N. (2002). New technologies and “cyber-queer” research. In D. Richardson & S. Seidman (Eds.), *Handbook of Lesbian and Gay Studies*. London: SAGE Publications.
- Walby, K. (2005). How Closed-Circuit Television Surveillance Organizes the Social: An Institutional Ethnography. *The Canadian Journal of Sociology*, 30(2), 189–214. <http://doi.org/10.1353/cjs.2005.0043>
- Wang, H., Chin, A., & Wang, H. (2011). Social Influence on Being a Pay User in Freemium-based Social Networks. *2011 IEEE 25th International Conference on Advanced Information Networking and Applications (AINA)*, 526–533. <http://doi.org/10.1109/AINA.2011.35>
- Warner, M. (1999). *The trouble with normal: Sex, politics, and the ethics of queer life*. New York: The Free Press.
- Waskul, D. D. (2002). The Naked Self: Being a Body in Televideo Cybersex. *Symbolic Interaction*, 25(2), 199–227. <http://doi.org/10.1525/si.2002.25.2.199>
- Weizman, E. (2002, April 24). Introduction to the politics of verticality. Retrieved April 18, 2014, from http://www.opendemocracy.net/ecology-politicsverticality/article_801.jsp

- Wesch, M. (2010). Youtube and you: Experiences of self-awareness in the context collapse of the recording webcam. *Explorations in Media Ecology*, 8.
- White House. (2012). Consumer data privacy in a networked world: A framework for protecting privacy and promoting innovation in the global digital economy. Washington, D.C. Retrieved from www.whitehouse.gov/sites/default/files/privacy-final.pdf
- Whitney, E. (2012, February 8). *Why Are Faggots So Afraid of Faggots? A Flaming New Anthology*. Retrieved December 3, 2014, from http://www.huffingtonpost.com/emerson-whitney/why-are-faggots-so-afraid-of-faggots_b_1256925.html
- Wiley, J. (1995). No Body is 'Doing it': Cybersexuality as a Postmodern Narrative. *Body & Society*, 1(1), 145–162. <http://doi.org/10.1177/1357034X95001001009>
- Williams, A. D., & Tapscott, D. (2006). *Wikinomics, How Mass Collaboration Changes Everything*. London: Penguin.
- Williams, R. (1961). *The Long Revolution*. London: Chatto & Windus.
- Winetrobe, H., Rice, E., Bauermeister, J., Petering, R., & Holloway, I. W. (2014). Associations of unprotected anal intercourse with Grindr-met partners among Grindr-using young men who have sex with men in Los Angeles. *AIDS Care*, 26(10), 1303–1308. <http://doi.org/10.1080/09540121.2014.911811>
- Winner, L. (1999). Do Artifacts Have Politics? In D. Wajcman & J. MacKenzie (Eds.), *The Social Shaping of Technology: How the Refrigerator Got its Hum*. London: Open University Press.
- Wittkower, D. E. (2014). Facebook and dramauthentic identity: A post-Goffmanian

- model of identity performance on SNS. *First Monday*, 19(4). Retrieved from <http://firstmonday.org/ojs/index.php/fm/article/view/4858/3875>
- Woo, J. (2013). Meet Grindr: How one app changed the way we connect. Self-published.
- Wortham, J. (2013, March 10). How Grindr is changing the way we connect. Retrieved March 10, 2013, from <http://bits.blogs.nytimes.com/2013/03/10/how-grindr-is-changing-the-way-we-all-connect/>
- Wright, L. (1997). The original Bears Mailing List: An interview with Steve Dyer. In L. Wright (Ed.), *The bear book: Readings in the history and evolution of a gay male subculture*. New York: Harrington Park Press.
- Wright, L., & Wehrle, D. (2001). The Bears Mailing List, part 2: Interview with Henry Mensch. In L. Wright (Ed.), *The bear book II: Further readings into the history and evolution of a gay male subculture*. New York: Harrington Park Press.
- Yiannopoulos, M. (2012, November 12). Did Grindr accidentally kill gay culture? Retrieved March 21, 2013, from <http://www.kernelmag.com/yiannopoulos/3256/did-grindr-accidentally-kill-gay-culture/>
- Yoo, D., Lake, M., Nilsen, T., Utter, M. E., & Alsdorf, R. (2013). Envisioning across generations: A multi-lifespan information system for international justice in Rwanda. *Science, Technology, & Human Values*.