



Human Data Commons Foundation

2017 Quantified Self Report Card: User Rights in the Age of Biometric Tracking

Project Overview	4
Report Card Layout and Navigation	5
Defining Privacy and Data Protection Terms	6
Standardizing Definitions: Personal Health Information	6
User Rights and Agency: "Opt-In" versus "Opt-Out" and "Condition of Use"	8
Users' Awareness and Choices: "Notice and Consent"	9
Quantified Self and Biometric Trackers: New Definitions and Challenges	11
Research Methodology	13
Experimental Design and Execution	14
Rating Criteria and Questions	16
1. Easy to Find Privacy Policy Link	16
2. Dedicated Privacy Contact	17
3. Information on Future Changes	17
4. Using Clear and Readable Language	17
5. Providing Means for Direct Contact	18
6. "Extra Credit" - Aspirational Information Provision	18
Ability to Import & Export Data	19
Ability to Delete Data at Will	19
Provision of Open API for Developers	19



Description of Technical Security Standards	19
Description of Storage Location and Practices	20
Defining and Categorizing Companies	20
Devices	22
Devices: Best Practices	24
Devices: Red Flags	25
Devices: Recommended Improvements	25
User Platforms	27
User Platforms: Best Practices	29
User Platforms: Red Flags	30
User Platforms: Recommended Improvements	30
Middleware Analytics	32
Middleware Analytics: Best Practices	34
Middleware Analytics: Red Flags	35
Middleware Analytics: Recommended Improvements	35
Conglomerates	36
Conglomerates: Best Practices	38
Conglomerates: Red Flags	39
Conglomerates: Recommended Improvements	39
Conclusion	41
Overall Industry Recommendations	41
Clarity, Navigability, and User Interfaces	42
Informed Consent and Readability	43
Comprehensive Coverage and Consumer Education	44
Research for Social Good	45
Methodological Assumptions and Limitations	46
Improvements for 2018 QS Report Card	47
Appendix A: Individual Company Reviews	50



In-Depth Reviews: Devices	50
In-Depth Reviews: User Platforms	56
In-Depth Reviews: Middleware Analytics	62
In-Depth Reviews: Conglomerates	66
Appendix B: Researcher Process Reflections	71
Reflections: Rochelle Fairfield	71
Concrete Research Process	71
What I Realized as I Went	71
What Should Be Done Differently in the Next Research Cycle	72
Reflections: Chelsea Palmer	73
Concrete Research Process	73
What I Realized as I Went	73
What Should Be Done Differently in the Next Research Cycle	74
Appendix C: Personal Post-Script	76
An Introspective Overview of the Cultural Relevance of Data Ethics by Rochelle Fairfield	76
Appendix D: Example Research Forms	79
Appendix E: Raw Data Spreadsheets	80
Appendix F: Works Cited	81



Project Overview

Since the rapid rise of computer and information technology, the privacy and security of user data has been a concern for many. With the mass societal adoption of the Internet, major companies and even governments now have access to accumulative databases of



all-encompassing user information aggregated over long periods of time. International regulatory organizations and activist nonprofits have proposed various frameworks to ensure that such databases are protected, rather than exploited.

When we choose to use these advanced and increasingly ubiquitous devices and services, we “opt in” to a new reality of constant monitoring which provides more complete “pictures” of our daily lives than we ever imagined possible. However, the companies which provide such devices and services often make it difficult for users to decide, or even understand, how their

data is being collected and utilized. Thus, we must quickly move to supplement the existing efforts of user rights watchdog organizations with a critical eye toward these new forms of data tracking and collection.

The [Human Data Commons Foundation](#) believes a new understanding of user rights in the Quantified Self field must start with an investigation of the user experience itself. A great deal of active research is underway to analyze the content of privacy policies with machine learning



technology, and legal and academic scholars are making consistent progress in establishing ethical norms for industry best practices. However, these approaches must be supplemented by a focus on the navigability of website and application interfaces “at first glance.” The average user cannot make choices about the privacy and security practices of a company if they can’t find or understand its documentation. Thus, we have undertaken a preliminary qualitative assessment of the perceived transparency and accessibility of such information in Quantified Self companies’ user-facing websites.

Report Card Layout and Navigation

The following report begins with [a brief overview of concepts and definitions found in the existing literature on privacy and user rights](#). It moves on to discuss the [research methodology and rating criteria for this “Report Card” evaluation](#), followed by the results of our analysis of companies within the Quantified Self field. These companies are [grouped within four categories \(Devices, User Platforms, Middleware Analytics, and Conglomerates\)](#), and the general themes and trends in each category are summarized. Finally, [overall industry recommendations for best practices](#) in user-facing privacy documentation are provided, as well as a [reflection on the limitations of this study](#) and [potential improvements for next year’s Report Card](#). The body of the Report Card is then supplemented with six appendices: [in-depth reviews](#) of each of the evaluated companies, [process reflections](#) from the two researchers, [a personal post-script connecting the research to the lived human experience](#), the [research forms](#) utilized in the study, [the raw data](#) that was analyzed, and the [works cited](#) within the Report Card.



Defining Privacy and Data Protection Terms

Whether in formal or informal contexts, any discussion about an abstract concept like “privacy” requires a collaboratively agreed-upon definition for the central terms used. In particular, the socially constructed notion of privacy has evolved alongside our representational technologies, moving from a dialogue about “the right to be left alone” during the initial rise of photography and mass media printing, to the contemporary dialogue about “the right to be forgotten” (Woo, 2006; Thompson, 2011; Vayena et al., 2013). Additionally, it has shifted from a right that solely benefits the individual to one that is crucial for the equality of larger society (Roessler & Mokrosinska, 2013).

As technology has become ubiquitous to most modern environments, and the majority of individuals experience significant benefits and convenience from its daily use, it has become useful to shift from a primarily restrictive to a critically reflexive framework regarding the boundaries between personal information and public life (Lederer, Dey, & Mankoff, 2002). However, the benefits users receive in exchange for, and because of, their provision of personal information cannot be assumed to wholly outweigh the risks that such collection poses to their personal rights. In order to further explore what those rights might be, we will first identify the terms to be used.

Standardizing Definitions: Personal Health Information

When it comes to legal documentation and formal policy-making, the most crucial definition to clarify in this discussion is the most basic unit of focus: personal information. There is an extensive body of literature navigating the factors which indicate a particular type of data could be sensitive or personally identifiable.

The majority of Quantified Self platforms and device manufacturers that we reviewed were based in the United States, so it is ideal to begin our definition of “privacy” from the accepted



legislation there. Data privacy policies in the United States are largely anchored in the content of the [Health Insurance Portability and Accountability Act \(HIPAA\)](#). This Act addresses the sensitivity of personal medical information, and holds industries and professionals responsible for carefully safeguarding it.

[Under HIPAA's definitions](#), "*protected health information*" includes two classes of data: *general health information*, and the subset of more highly sensitive *personally identifiable information*, containing elements which could allow that information to be linked to a specific individual. Even outside of the targeted medical field, major technology platforms are including personal health information as a category in [their internal definitions of sensitive personal information](#).

Various degrees of sensitivity and identifiability can be further distinguished, from information that is potentially "linkable" to someone's identity by investigating combined traces of other factors, to data which is directly "linked" to someone's identity by clear markers such as their name or social security number. The guidelines of HIPAA only apply to personal information that can be identified in the sense that it is reasonably "linked" or "linkable" by these definitions, and [has not been anonymized or de-identified](#). Most relevant to our research, biometric records are considered a positive linked correlate to an individual's identity, rather than merely a potentially linkable attribute (McCallister, Grance, & Scarfone, 2010).

Outside of the United States, there are diverse national approaches to defining and regulating forms of protected health information. The designating characteristics of "sensitive" personal information within these various regulatory structures generally agrees with those established under HIPAA, [though EU legislature uses the term "personal data" in place of "personally identifiable information."](#) The primary distinction between US privacy laws and those in Canada, Europe, Australia, and elsewhere is the seeming "default state" incurred by policy design: the United States is known for a largely *permissive* approach to data collection, whereas the European Union is known for a more *protective* approach to such collection (Thierer, 2015;



Weiss & Archick, 2016). This suggests a fundamental divide in attitudes regarding users' rights and agency to make decisions in these matters.

User Rights and Agency: “Opt-In” versus “Opt-Out” and “Condition of Use”

The permissive “free market” data collection policies of the United States are based upon the assumption that restrictive regulations in the private industry would not only stifle technological innovation, they would also deprive users of the agency to make their own choices (Thierer, 2015). In contrast, more restrictive data collection policies in the EU are driven by the assumption that without firm government regulation, private companies would find ways to completely strip the user of any right to choose whatsoever (Leibenger et al., 2016). Of course, we can find strengths and weaknesses in each argument, and the truth probably lies somewhere in between.

Interestingly enough, these two ideological perspectives correspond closely with two concrete models of privacy notification. These are referred to as “opt-in” and “opt-out” style practices (Vayena et al., 2013). With “opt-in” practices, common under EU regulation, the user must explicitly consent to any specific conditions of data collection or use outside of the purpose of a given service (Leibenger et al., 2016). In contrast, the “opt-out” practices that are widespread in the US tend to grant the service provider broad, unrestricted access to the users' data for any legally acceptable use as soon as they sign up. In many cases, this “opt-out” structure is further restrictive of users rights, by offering very few legitimate options to “opt-out” of anything while maintaining an active account: this functions essentially as a “condition of use” situation, in which the only two true choices are to grant broad consent to the terms as they're laid out or to not use the services (Vayena et al., 2013).

The lived experience of the “opt-out” method is viscerally memorable for anyone who has spent what seems like hours of their life going through a new account's settings and “un-selecting” a flurry of auto-subscribed emails and permissions. There is also no clear method of ensuring



that the “opt-out” process is fully successful, whether from the possibility of continued covert data collection, or these options just seeming to “reset” to tracking the users’ data after service updates. In sum, the “opt-out” structure puts the power of setting the terms of engagement in the service provider’s hands, and the responsibility for contesting those terms in the user’s hands.

Under “opt-in” enforced practices, the user would only have to dedicate time to clicking through every settings menu if they wanted to grant more access and information to the service provider, rather than less. The “opt-in” structure does not necessarily put the power to set terms in the user’s hands, but rather seems to ensure a more neutral starting point, in which neither party is presumed to agree to anything more than the exchange of information and services that is at the core of their shared business.


Users’ Awareness and Choices: “Notice and Consent”

Under current industry regulations, the above operational description of “opt-out” settings is enacted via the norm of “notice and consent” privacy practices. “Notice and consent” is part of the “Fair Information Practices” introduced by the FTC (Marotta-Wurgler, 2015; Sathyendra et al., 2016). Upon the user’s first access to the company’s site, this practice informs them of the overall scope of the company’s general privacy practices. “Notice and consent” then requires the user’s affirmative acknowledgement that this information has been received (and in an ideal world, understood).

“Notice and consent” has been shown to have little effectiveness in actively and thoroughly informing users of their rights (Marotta-Wurgler, 2015). Rather than sticking with the problematic binary notion of user consent in the “notice and consent” model, researchers are proposing that we move toward a definition of consent as a highly contextual choice, for which there should be adaptable options given different situations of data collection and use (Vayena et al., 2013). Since the user’s contribution of data to a platform is an inherently participatory



process, a more complex and adaptable interface than mere “notice and consent” would reflect and respect that agency (Woo, 2006).



Quantified Self and Biometric Trackers: New Definitions and Challenges

The enthusiastic Quantified Self movement is emblematic of the swift changes that Big Data collection and analytics have brought to our society, and its participants are poised to reap both



the rewards and the potential risks of this cutting-edge technology. Whereas [personal health information was once bounded by paper records controlled by medical professionals](#), QS enthusiasts can now collect and analyze inclusive profiles of their health and well-being, with increasingly sophisticated tools (Leibenger et al., 2016).

Though the industry is most popularly known for activity and fitness trackers, diverse devices are entering the market to track more subtle biological characteristics like ketone levels and brainwave activity (Haddadi & Brown, 2014). Additionally, the adoption of these devices and platforms is expanding beyond QS communities into widespread populations, [with some professional medical perspectives claiming that such device use encourages long-term commitment to healthy regimens](#) (Piwek et al. 2016).

Though it is widely known and implemented, the United State's HIPAA privacy framework has a significant gap for the Quantified Self community: [its policies are only enforced in the practices of health care providers and health care plan systems](#), and [do not apply to the rising industry of consumer health trackers and wearables](#). Therefore, while HIPAA contributes to our establishment of privacy field definitions, [these Quantified Self products and services are produced and maintained by "non-covered entities" within the American HIPAA framework](#), and thus oversight has been largely left to consumer market regulations, (Acquisti et al., 2014). In the United States, this responsibility has largely fallen to [the Federal Trade Commission](#), which sporadically enforces strong policies against major industry conglomerates due to data breaches caused by perceived failure to follow best security practices (Mehlman, 2015).

The data brokerage industry, buying massive reserves of information from customer-facing companies, is quickly becoming one of the most profitable sectors in the technology industry (Mehlman, 2015). So far, the benefits of this big data boom have been completely asymmetric, with the users contributing this information being pacified with the false impression that the



services that would already be needed for the collection are somehow a fair financial trade-off (Andrejevic, 2013).

One risk incurred by embracing biometric and GPS tracking is the consistently demonstrated capability to [reliably predict someone's future location based upon a record of their previous movements](#). On another front, there is the pessimistic possibility that, given access to health tracker databases, private sector health insurance providers will begin to discriminate against clients based on this information, in a similar fashion to denying coverage based on pre-existing conditions (Lupton, 2017).

Research has made it abundantly clear that users of wearable technology are aware of these risks, and [many of them express proactive concern about shaping appropriate regulations as this field expands](#). Additionally, the line between mainstream devices and targeted activity tracking wearables grows increasingly blurred, as [major device manufacturers increasingly tailor smartphone hardware to enable built-in Quantified Self capacities](#).



Research Methodology

We approached this initial Report Card with the simplest framework possible. Given the relatively recent advent of Quantified Self tracking technology, we focused on laying the foundation for a stable but flexible long-term monitoring structure in this field. Thus, some of the more visionary elements of the Quantified Self Report Card were reserved for the 2018 edition and beyond, as we navigated unforeseen challenges in our first round of research.

Our initial intention was a thorough review of all the elements of privacy and security within each company's practices. As the research progressed, it became clear that the most relevant exploration for the 2017 Report Card would involve scrutiny of the ease of access to each company's user-facing information. More specifically, we aimed to evaluate how easy or difficult it seemed for a website visitor to navigate the company's website interface and track down answers to any questions they might have about data privacy and security practices.

Originally,, the idea of engaging the Quantified Self and technology communities in a widespread "crowdsourcing" campaign was a driving aspiration. However, as many researchers have highlighted, crowdsourced evaluation of privacy policy content is an excessively time-intensive process (a reality that was further confirmed by the time our own researchers devoted to each company's evaluation). Thus, the vision of this widespread, pseudonymous participation was not included in this year's Report Card edition.

However, crowd review and feedback on the experimental design and theoretical focus was achieved throughout the research process via three public discussion groups, conducted once every six to eight weeks, held onsite at [dctrl community commons](#) in Vancouver, BC. This allowed the researchers to revisit their assumptions about users' priorities throughout the experimental design process, and adjust their methodological focus as necessary.



After reframing our methodological review forms based on community feedback, we came to find the unique value add that we could provide within the existing body of research on user-facing privacy practices. By approaching the interface of these sites with reflexive self-awareness, rather than attempting to establish some sort of rigid and objective criteria, our inquiry was indicative of our actual lived experiences as Internet users.

The most careful and granular academic analysis of privacy policies' text is rendered useless if no one can track it down on a company's website. Similarly, while machine learning can be trained to recognize certain words or terms as vague, the subjective judgment of a human reader is still able to follow the more subtle overall degree of perceived clarity or confusion present within a piece of text (Sathyendra et al., 2016). In short, the very same aspects of our methodology that would negate its validity in traditional research paradigms are actually its strengths in terms of a more humanist and post-structural evaluation.

Experimental Design and Execution

The initial collection of companies and devices for potential review was led by the Human Data Commons Foundation's Board Chair, Scott Nelson, in Fall 2016. A cohesive spreadsheet was created, with entries added over the span of a few months. To bolster the effective coverage of the Quantified Self field, the Foundation [reached out to the Reddit community seeking additional suggestions for this spreadsheet](#) (the lack of responses to this reflected a consistent discovery: the challenges of engaging widespread crowd-sourced participation). Less than ten companies from the initial list were dropped from the final results upon initial review, due to being a poorer fit for the Quantified Self focus than predicted, or due to the company not being far enough along in development for a fair assessment. In early Spring 2017, the list was "closed" and internal research began.



The core execution of the evaluative research was conducted by two Human Data Commons Foundation researchers (Rochelle Fairfield and Chelsea Palmer), who separately visited a list of about fifty websites and filled out a standardized research form for each site. The two researchers' forms, which are available for review as [Appendix D](#) to this Report Card, were each tailored to slightly different ends. Rochelle's form was focused on the immediate user interface and experience upon visiting the site. Chelsea's form was focused on a deep-dive through the entirety of the websites' available information, looking to "fill in the blanks" with less immediately available answers to the research questions..

Whereas the preliminary experimental design attempted to pursue rigid objectivity, this approach was quickly adjusted to focus on the researchers' individual perceptions. For example, while we could not reliably find out each company's technical security standards, we could certainly comment on whether or not it was possible to find any information on those standards in less than ten minutes of searching.

The process of review elapsed over a period of less than one month, with a concrete attempt made to allot a similar amount of time to each company's site. Upon completion of these forms, Chelsea Palmer handled the data export, cleaning, and organization. The Google Forms results were initially exported into raw data spreadsheets, which was then processed manually, due to the subtlety of many of the open-ended questions and the slight differences between the two researchers' targeted forms. Three iterations of data organization unfolded during this process (the full set of raw data spreadsheets is provided within [Appendix E](#)), as these increasingly organized evaluations were utilized to determine the most effective structure for comparative rating and presentation.

Though the original research forms inquired about both the Privacy Policy and Terms of Service agreements, in processing the data it became clear that the Terms of Service agreements were almost always in rigid and dense legal language. This is understandable given their inherent



purpose: to protect against every potential legal vulnerability for the company that produces them. However, the Privacy Policy of a company should ideally exist to inform the end user of their rights, rather than serving the company alone. Therefore, the focus during final data analysis was primarily on the Privacy Policy as a means to inform and ideally empower the end user of a company's services.

The data analysis process was primarily focused on highlighting potential gaps in the existing methodology, in order to strengthen future editions of the Report Card. Areas in which the two researchers provided conflicting answers, even to subjective questions, were considered the most fruitful for understanding both weaknesses in our research methodology and ambiguity in the interface of the sites themselves. In these cases of conflicting researcher results, answers about interface clarity and navigability were deferred to Rochelle's feedback, whereas answers about the ultimate availability of information on technical topics were deferred to Chelsea's feedback.

Rating Criteria and Questions

After processing the data, the longer research forms were distilled into five core questions for comparative "yes/no" rating, as well as five questions for aspirational best practices on in-depth disclosure, which would add up to a single "extra credit" point, for an overall possibility of six points for each company's rating. The criteria for these rating questions is more clearly explained as follows:

1. Easy to Find Privacy Policy Link

In general, standards for best practice data privacy disclosure require a link to the company's Privacy Policy to be clearly displayed and accessible on each page (ideally fixed within site's the header or footer). This criterion's binary point (either "one" or "zero") was determined to apply solely to the Privacy Policy link, for the reasons described above, but the few cases in which the



Privacy Policy link was provided, but the Terms of Service link was absent, are noted within the detailed company reviews of [Appendix A](#).

2. Dedicated Privacy Contact

Another element of best practices requires a dedicated privacy contact within the company. Ideally, this should be a designated and named Chief Privacy Officer, but as the adoption of this approach is still very limited, for the purposes of this Report Card this binary point (either “one” or “zero”) was allotted to any company that designated a specific contact solely for the purposes of privacy inquiries, even if it was just a “privacy@” email address.

3. Information on Future Changes

The user’s relationship with a company’s Privacy Policy does not end with their initial agreement to its terms. It is crucial that companies establish if, and how, users will be informed of changes to their Privacy Policies. A point (either “one” or “zero”) was allotted for merely disclosing this information, whereas commendations are given in [Appendix A](#) for those companies that pledged to proactively inform users of changes, rather than leaving the user solely responsible to check back for potential updates.

4. Using Clear and Readable Language

This is obviously the most subjective rating criterion within our experimental design, but it is also the most important for the spirit and intent of the Report Card project. In order to mitigate the inevitable biases that come with personal perspective, effort was exercised to look in depth at the overall feedback from both researchers for each company’s policies.

Each researcher’s form had two yes/no questions about readability: whether they initially found the policies themselves to be in easily readable language, and whether after full review of the site they felt the company had exercised effort to present its information in an approachable and readable fashion. Additionally, the researchers provided open-ended notes qualifying these



binary responses within the comments sections. During data analysis, each company was scored based on an evaluation of all of this feedback brought together. When there was some conflict or unsureness in the researchers' perceptions, this was noted in the detailed company reviews in [Appendix A](#).

5. Providing Means for Direct Contact

There has been an increasing shift over the past decade from avenues to direct human contact to static "Frequently Asked Questions" pages and comment submission forms. This is understandable in terms of companies' resource allocation, but it can lead to significant frustration and time wasted for users. By providing multiple avenues for users to submit their questions or feedback, companies take on proactive responsibility to directly respond to such inquiries. Thus, aggregate points (each counting as 0.25 out of one full point) were afforded for four methods of potential outreach: email address, phone number, physical address, and static comment submission form.

6. "Extra Credit" - Aspirational Information Provision

Five questions from the researchers' review forms were determined to represent more aspirational practices of information disclosure, rather than core requirements for the basic Privacy Policy expectations. Additionally, the researchers' responses to these questions was acknowledged as comparatively imprecise in terms of objective answers, as acknowledged in the [review of our methodology limitations](#).

While we could address our ability to find surface information about the following criteria, as subjective reviewers we would inevitably miss some of the answers within the documents provided. To account for this, each of the following criterion was measured as an aggregative benefit, based on the ease of uncovering direct answers to such questions within the surface documentation provided. Each question was counted as 0.2 points, for a total availability of one



full point for the final Report Card rating. The results tables for these “extra credit points can be found within [Appendix A](#).

a. Ability to Import & Export Data

This fractional point was awarded to any company that explicitly highlighted the users’ ability to import data from other services or devices into the company’s platform, and/or export their aggregated data from that company’s control into external services. Instances in which only import or export were addressed received the fractional point, but the incomplete information was acknowledged in the detailed reviews of [Appendix A](#).

b. Ability to Delete Data at Will

This fractional point was awarded to any company that explicitly highlighted the users’ ability to successfully request deletion of their data from the company’s service or platform. In this criterion rating, no distinction was made based upon the actual availability of such a deletion option, merely on the company’s transparent disclosure around deletion practices. More detailed discussion of these practices is included in [Appendix A](#).

c. Provision of Open API for Developers

This fractional point was awarded to any company that explicitly highlighted the availability of a publicly accessible API for technical developers to design integrative services for the company’s platform. Though we designate the criterion as an “open” API, it was often difficult to determine whether permissioned APIs were considered fully open. Therefore the point was awarded based on links to the API portal, but future Report Card editions will delve into a more technical review of whether such portals are truly “open,” or maintained under full proprietary control by the companies.

d. Description of Technical Security Standards

This fractional point was awarded to any company that provided information on the technical security standards of the company’s data transmission and credential verification. There was a



significant variation between the two researchers' perceptions of whether companies effectively disclose technical standards, so the text of these disclosures was directly copied into the comprehensive result spreadsheets. Many companies which provided such information did so in a rather general fashion, so the degree of specificity in these disclosures is discussed in more detail in [Appendix A](#).

e. Description of Storage Location and Practices

Parallel to the description of technical standards, this fractional point was awarded to any company that provided information on where and how user data is stored. This varied from descriptions of the geographical location of servers, to disclosure about control over those servers, and local storage of data on user devices. Again, the degree of specificity in these disclosures is discussed in more detail in [Appendix A](#).

Defining and Categorizing Companies

One of the central distinctions that became clear as we surveyed user facing privacy practices in this field reflected the saying “not all organizations are created equal.” It would be disingenuous to compare the progress and resources of a device manufacturing startup launched within the past year with those of a decades-old technology conglomerate. Similarly, there are distinctive differences in the presumed responsibilities of different company categories: a consumer-facing platform should provide more focus on informed user consent than a “middleware” provider which distributes technologies for others to build their services upon.

Of course, drawing such defining lines is imperfect, as some companies could be subjectively defined under more than one category, but the categories were based upon the perceived primary function of each company. Additionally, general patterns emerged in the research data outcomes which enabled further clarity in defining these comparative categories.



The category “[Devices](#)” refers to the field we expected to survey from the very beginning: companies which primarily manufacturer equipment and tools used for self-tracking.

The category “[User Platforms](#)” refers to applications and websites which primarily provide user-facing platform services, from analytics for raw biometric data from wearable devices to interfaces for self-reported qualitative data.

The category “[Middleware Analytics](#)” refers to companies which offer analytics and algorithms for other companies to build their user-facing services upon.

The category “[Conglomerates](#)” refers to companies which were already established in the technology industry before venturing into the Quantified Self field, and which still provide products and services outside of biometric and self-tracking devices and platforms.

The industry overview to follow will touch upon the notable strengths and weaknesses within each overall category of organization as defined above, with an eye to highlighting examples of best practices. In order to focus on the larger trends of these categories, the intention choice has been made to avoid referring to specific companies by name within the following section. [Appendix A](#) provides more thorough qualitative reviews of each specific company.



Devices

Though review of this category is the easiest to standardize, these companies face the most challenges in terms of establishing norms for transparent, informed user consent. As the “ground floor” of data collection, the actual biometric trackers used in the Quantified Self field serve as the frontline for defining the relationship between users and their data.

Score Legend:	Colour:				
	Rating:	Below Average	Average	Above Average	Excellent
	Points (out of 6):	0 - 2	2 - 3.5	3.5 - 5	5 - 6

Company	Privacy Policy Link	Dedicated Privacy Contact	Future Changes	Attempts at Readable Policy	Points of Direct Contact (out of 4)	Extra Credit	Final Score
Fitbit	1	1	1	1	0.75	0.8	5.55
Muse	1	1	1	1	0.75	0.6	5.35
SigmaSport	1	1	1	1	1	0.4	5.4
Withings	1	1	1	1	0.25	1	5.25
Mio	1	1	1	1	0.75	0.4	5.15
Misfit	1	1	1	1	0.75	0.4	5.15
Nervana	1	1	1	1	1	0	5
Suunto	1	1	1	1	0.5	0.4	4.9
Bloomlife	1	1	1	1	0.5	0.2	4.7
Bellabeat	1	0	1	1	1	0.4	4.4
Wahoo	1	0	1	1	1	0.2	4.2
Blocks	1	0	1	1	0.25	0.6	3.85
Emotiv	1	0	1	1	0.25	0.6	3.85
Sidly	1	0	1	0	1	0.2	3.2
Athos	1	0	0	1	0.75	0.2	2.95
Soleus	1	0	1	0	0.75	0.2	2.95
Polar	1	0	1	0	0.5	0.2	2.7
Ketonix	0	0	0	0	0.75	0.2	0.95
Tic	0	0	0	0	0.5	0.2	0.7
iWinks	0	0	0	0	0.25	0	0.25




Privacy Policy Link




The majority of these companies provided a fixed link to the Privacy Policy in the footer of the website. The three companies that did not do so were the lowest scoring devices within the Devices category.

Privacy Contact



In contrast, within the Devices category, the lowest scoring criteria was the provision of a dedicated privacy contact. This can perhaps be attributed to the relatively small size of many of these companies.

Indicate Changes




The majority of companies within the Devices category provided information about how future changes to the policies would be indicated.

Readable Policy



The researchers perceived that about two thirds of the companies within the Devices category had made an attempt to use easily readable language within their Privacy Policy descriptions, rather than legalese.



The majority of companies in the Devices category provided more than one point of direct contact for general inquiries (email, physical address, phone number, contact form). This was one of the greatest strengths of the Devices category overall.

Devices: Best Practices

The majority of devices met the four core criteria for the Report Card, with particularly strong adherence to the requirement for fixed Privacy Policy links on their websites, and an indication within the Policy of how future changes would be indicated. Additionally, the researchers shared



a positive subjective evaluation of the majority of these companies' attempts at providing user-friendly Privacy Policies which used simple and approachable language. The majority of companies were transparent about users' ability to delete their data and/or accounts.

Devices: Red Flags

The three lowest scoring companies did not meet any of the four core criteria for the Report Card, lacking fixed website links to the Privacy Policy, a dedicated privacy contact, information about future change indications, and apparent attempts at an easily readable Privacy Policy, with one company lacking a Privacy Policy together.

Additionally, the criteria for the "extra credit" points were often initially difficult to uncover: in particular, the users' ability to import or export data was rarely addressed. This category had the most frequent occurrence of responses of "unsure," rather than "yes" or "no," from Rochelle's user-centric evaluation. Many of the companies whose Privacy Policies were reported as difficult to navigate had significant problems in this regard. For companies in other categories, poor navigability usually referred to dense language and excessive length, whereas within the Devices category it ranged from apparent translation barriers to sections or documents that were misplaced or missing entirely.

Devices: Recommended Improvements

The production of actual measurement devices stands at the frontline of Quantified Self industry practices: as opposed to behind-the-scenes analytics or self-report platforms, device manufacturers collect the raw data of users' unconscious movement and activity. Due to this granular, all-encompassing collection, is crucial that device retailers clearly and comprehensively inform users about what information is being collected and what it is being used for, and to provide flexible options to control the import, export, and deletion of this information.



Device manufacturers have to both address the intense “linkability” present within their collected data reserves, which often simultaneously track multiple aspects of the user’s activity, while grappling with the challenge that increasingly small and unobtrusive device interfaces present for the “informed notice” approach to disclosure (Lederer, Dey, & Mankoff, 2002). If manufacturers embrace the challenge of this field’s relatively new foundation, they can set cutting-edge best practices for long-established technology giants to turn to for inspiration– or guilt– to change stagnant and opaque precedents.



User Platforms

The category of User Platforms is less of a distinct definition than the others, as a few different types of interfaces and platforms fall under this title. However, in general, these platform companies are collecting and handling “second hand” biometric data imported from other devices, or storing and providing feedback on self-reported, often qualitative, information submitted by users.

Score Legend:	Colour:				
	Rating:	Below Average	Average	Above Average	Excellent
	Points (out of 6):	0 - 2	2 - 3.5	3.5 - 5	5 - 6

Company	Privacy Policy Link	Dedicated Privacy Contact	Future Changes	Attempts at Readable Policy	Points of Direct Contact (out of 4)	Extra Credit	Final Score
Clue	1	1	1	1	0.5	0.8	5.3
X2AI	1	1	1	1	0.75	0.4	5.15
Beeminder	1	0	1	1	1	0.6	4.6
Dacadoo	1	0	1	1	1	0.6	4.6
My Fitness Pal	1	1	1	1	0.25	0.2	4.45
Fitness Syncer	1	0	1	1	1	0.2	4.2
Predict BGL	1	1	0	0	1	1	4
HeartMath	1	1	1	0	0.75	0.2	3.95
LibreView	1	1	1	0	0.5	0.4	3.9
BioBeats	1	1	1	0	0.25	0.6	3.85
SAMapp	1	1	0	1	0.25	0.4	3.65
Training Peaks	1	0	1	0	0.5	0.8	3.3
Nutrino	1	0	1	0	1	0	3
Pacifica	1	0	1	0	0.25	0.4	2.85
Nutra Hacker	1	0	0	1	0.25	0.4	2.65




Privacy Policy Link



All of the companies in the User Platform category provided a fixed link to the Privacy Policy across their website.

Privacy Contact



A little more than half of these companies provided a direct point of contact for privacy inquiries.

Indicate Changes




The majority of these companies clarified how they will indicate future changes to the privacy policy to their users.

Readable Policy



A little more than half of these companies had privacy policy documentation that the researchers perceived to be easily readable for an average user.



All of the user platform categories provided at least one point of contact, with about one third of them providing all four possible points of contact.

User Platforms: Best Practices

Most impressively, all User Platform companies had fixed links to the Privacy Policy within the header or footer of their websites. Additionally, the majority of these companies disclosed their approaches to informing users about future Privacy Policy changes. All companies provided at least one avenue to directly contact a human being, and the majority of them also gave some indication of their technical security standards, although they were often generalized.



User Platforms: Red Flags

Reviewers evaluated that only about half the companies seemed to make an effort to provide easily readable and direct information within their Privacy Policies, with many of them instead relying upon seemingly boilerplate Policies comprised of dense legal language.

It was also difficult to uncover information from most of these companies about the options for users to import, export and delete their data. This is concerning, as it is especially crucial for User Platforms, which often act as a dashboard for aggregated data from various sources. If the reserves of users' data is siloed within these proprietary platforms, they are left with little choice to switch to other platforms without "starting from scratch" in terms of data collection. Finally, although all companies provided contact information, in some cases it was not particularly easy to track down, and could be displayed more prominently.

User Platforms: Recommended Improvements

This is the category where it is most crucial to lay down stronger definitions and distinct terminology. Many platforms are beginning to combine self-report data with that funnelled directly from other services and devices, aggregating huge amounts of data about a single user in one place. They need more than anyone to break down the difference between these various types of information, and transparently disclose how they store it and manage the inevitable attack vectors of interoperable exchange. Unique to the user platforms is the issue that the stored data often faces a number of different parties: the users themselves, the service providers, advertising agents, and other users in a public-facing capacity. The disclosure of data practices should clearly break down the different levels of sharing and exposure for these different actors.

Similarly, since user platforms often incorporate an element of social media interaction, the comparative sensitivity created by this combination and display of one's info needs to be



addressed adequately - this returns to the concept that privacy is not merely about retaining information for sole individual use, but navigating the acceptable social boundaries of sharing information (Roessler & Mokrosinska, 2013).



Middleware Analytics


In one respect, middleware companies have a slightly less urgent need for transparent and comprehensible documentation of privacy rights, as they do not provide direct services to end users. In another respect, they have an even greater responsibility to set precedents for the companies to whom they sell their technology. Unlike the average user, who currently may face a “take it or leave it” attitude from such service providers, Middleware companies hold the bargaining chip of innovative and cutting edge tools which these service providers need in order to operate.

Score Legend:	Colour:				
	Rating:	Below Average	Average	Above Average	Excellent
	Points (out of 6):	0 - 2	2 - 3.5	3.5 - 5	5 - 6

Company	Privacy Policy Link	Dedicated Privacy Contact	Future Changes	Attempts at Readable Policy	Points of Direct Contact (out of 4)	Extra Credit	Final Score
Validic	1	1	1	1	1	0.6	5.6
Eyeris	1	1	1	1	1	0.2	5.2
Affectiva	1	1	1	0	0.75	0.2	3.95
Qusp	1	0	1	1	0.75	0	3.75
Imotions	1	0	0	1	0.75	0.6	3.35
Beyond Verbal	1	0	1	0	1	0.2	3.2
Advanced Brain Monitoring	1	0	0	1	0.75	0.4	3.15
Fatigue Science	1	0	0	1	0.75	0.2	2.95
Human API	1	0	1	0	0.5	0.4	2.9
Sensaura	0	0	0	0	0.75	0.4	1.15



Privacy Policy Link



The majority of companies in the Middleware Analytics category - all except for one - provided a fixed link to their Privacy Policy throughout the website.

Privacy Contact



The majority of these companies did not provide a dedicated point of contact for privacy inquiries, making this their weakest criterion - only three of them did so.

Indicate Changes




More than half of these companies included information about how the user would be informed of future privacy policy changes.

Readable Policy



More than half of these companies provided policies which the researchers perceived to be easily readable - this is particularly impressive given that the Middleware companies are not generally consumer-facing, and so more technical documentation would not be surprising.



The Middleware Analytics category had the best practices for direct contact, with all of the companies providing at least two such methods, and the majority providing three or more.

Middleware Analytics: Best Practices

All but one of the Middleware companies displayed appropriate fixed links to their Privacy Policies across their websites. Additionally, all but two of the companies gave information on how future changes would be indicated, although the majority of them placed the responsibility for tracking these changes on the user proactively checking the website.

The researchers perceived that about half of these websites seemed to have made an effort to use easily readable language within their documentation, and provided multiple avenues for



direct contact regarding general inquiries. Finally, the majority also provided at least some information about the technical standards they used, which is particularly important for the Middleware category, as they exist as the “bottom layer” for other companies’ operations.

Middleware Analytics: Red Flags

All but three of the companies lacked a dedicated privacy contact within their documentation, and as described above, only half of the companies were perceived by the researchers to have made efforts for easily readable documentation. There was also limited or unclear information about data import, export, and deletion, although this is understandable due to the fact that Middleware companies are not directly user-facing, and may leave it up to platforms to design their own policies in this regard.

Middleware Analytics: Recommended Improvements

Right now many middleware companies do not take a public stance one way or another on users’ digital rights, because they haven’t had to. Implementing policies and agreements which hold their company clients responsible down the chain to their end users would be a massive step forward to incentivize widespread adoption of these standards. Being in the strongest position to influence user-facing companies, they can encourage significant movements toward social good research without much effort on their own part, as they can wield the leverage of controlling and producing the cutting edge technology that the Quantified Self field thrives upon.



Conglomerates

The greatest strength of these conglomerates is their comparatively vast financial and legal resources, and general visibility. However, their size often results in a significant barrier to access in terms of contact with real human beings for inquiries or concerns. The greatest concern within this category was the lack of options for directly contacting a human being with inquiries. However, the Conglomerate companies had the highest rate of achievement for “extra credit” information disclosure, although these disclosures were often rather vague and did not address specific practices.

Score Legend:	Colour:				
	Rating:	Below Average	Average	Above Average	Excellent
	Points (out of 6):	0 - 2	2 - 3.5	3.5 - 5	5 - 6

Company	Privacy Policy Link	Dedicated Privacy Contact	Future Changes	Attempts at Readable Policy	Points of Direct Contact (out of 4)	Extra Credit	Final Score
Under Armour	1	1	1	1	0.5	0.6	5.1
Motorola	1	1	1	1	0.25	0.6	4.85
Telus	1	1	1	1	0.5	0.2	4.7
Apple	1	0	1	1	0.5	0.8	4.3
TomTom	1	1	1	1	0	0.2	4.2
Garmin	1	1	1	0	0	0.8	3.8
Google	1	0	1	1	0	0.8	3.8
Samsung	1	0	1	1	0.75	0	3.75
Microsoft	1	0	1	1	0.25	0.2	3.45
Timex	1	0	1	0	0.5	0	2.5




Privacy Policy Link



All of the companies in the Conglomerate category provided a fixed privacy policy link across their websites.

Privacy Contact



Exactly half of these companies provided a dedicated point of contact for privacy inquiries - this is particularly concerning given the resources they have at their disposal.

Indicate Changes




All of the companies included information on how the user would be notified of future changes to the policy.

Readable Policy



The majority of these companies provided documentation that the researchers felt to be easily readable by the average user.



The Conglomerates category had the weakest showing in providing direct points of contact, with nearly one third of the companies providing no means for contact at all, and only one company providing three points of contact.

Conglomerates: Best Practices

Most impressively, all Conglomerate companies have an appropriate fixed link to the Privacy Policy in the website's header or footer. In their documentation, all of these companies disclosed the means by which future Policy changes would be indicated. They also provided the most thorough documentation on privacy practices, and often shared dedicated overall statements about aspirational ethical practices.



The majority of these companies also provide explicit educational content addressing basic security protocols, informing users about how to best protect their accounts' security on their end. This is a crucial practice to address the "weakest link" risk within cybersecurity, as even the best company practices cannot compensate for weak passwords or capitulation to "phishing" attempts on the user side.

Conglomerates: Red Flags

Surprisingly, the majority of Conglomerate companies lacked a dedicated officer or email address for direct privacy inquiries. This was particularly remarkable given the thorough attention paid to overall education within "privacy portals". Additionally, the Conglomerate category was the weakest in terms of providing points of human contact in general, with three companies providing no means of contact at all. Finally, though most of these companies scored very high on this Report Card's criteria, the Policies provided often applied to the operations of the companies as a whole, with an absence of attention to specific Quantified Self devices and the uniquely sensitive personal health data which they collect.

Conglomerates: Recommended Improvements

Given the massive resources available, conglomerates should be held to a very high standard for security and privacy practices and disclosure. The huge scale of their physical storage operations and ubiquity of their services' public presence should prompt them to disclose much more concrete technical information, in order to appear as leaders in the field and relatively benevolent educators for consumers.

For similar reasons, these companies need to provide specific policies and easily accessible information focused on each of their individual services and devices, rather than a generalized universal company policy (which often seems to exist solely to limit their liability). The Conglomerate category had the most frequent instances of providing massive, detailed policies



that ultimately communicated very little useful information about their Quantified Self devices and services.

Frankly, it is completely unacceptable to have no means of direct contact for inquiries, and it's something that people regularly bemoan when a problem arises with one of these companies. If startup companies can have friendly email exchanges with less than 48 hours turnaround, there is no excuse for the major players to escape this. Additionally, every single company that has the massive reach and power to be designated a "conglomerate" has the ethical responsibility to pave the way for alliances with social good research. In the current age, the concept of corporate social responsibility needs to extend beyond one-time financial donations for tax write-offs, and begin to incorporate the value of personal data as a contribution to the greater good of society.



The Quantified Self consumer industry is at an exciting point in its development: what was once the realm of inquisitive hobbyists is gradually spreading to mass adoption and familiarity. This positions companies that collect biometric and self-tracking data to blaze forward and set new standards on a trail that's often been muddied and misleading: that of ethics in data collection.

Implementing “privacy by design” from the outset, the Quantified Self sector can ameliorate long-standing ethical conflicts within the technology sector (Thierer, 2015). Privacy by design moves privacy practices toward a user-centric, rather than data-centric, design of policies, without stifling the strong research opportunities enabled by big data collection (Adams & Sasse, 2001). Some of the improvements that can supplement this progress are grouped below, focusing on meeting the core standards of this Report’s methodology, clarifying language for informed user consent, and working toward relationships for social good research.

Clarity, Navigability, and User Interfaces

It will be a relatively straightforward process for companies to ensure compliance with four of the five core evaluation criteria: provision of fixed links to the Privacy Policy, a contact point for privacy concerns, disclosure of how future Policy changes will be indicated, and multiple avenues for users to contact a human being with general inquiries.

The issue of notifying users of changes to privacy policies is a challenging one, as “notification fatigue” can paradoxically lead users to get annoyed when companies enact best practices by emailing about updates (CITE). Nonetheless, these changes must be indicated in a manner that does not put the primary burden upon the user, as any changes to the policies will be considered as implicitly accepted by those who already use products or services.

Many companies that do reference future notifications of changes use the vague terminology “material changes” to indicate the threshold of perceived importance at which they will email customers directly. If this term can be more specifically defined, and standardized across



industry usage, then this is indeed an ideal approach to the issue— notifying users when there are changes to practices which could realistically impact their normal use or the overall treatment of their data.

The excessive length and overly complex navigation structures of many existing privacy policies serves as a significant barrier for users' informed consent (Schaub et al., 2015). Utilizing unique visual interfaces to summarize this information, such as icon indicators or rating matrices, could significantly increase user engagement in privacy-related decisionmaking (Alohaly & Takabi, 2016). Currently, front-end developers and copywriters are given few resources to navigate this design problem, basing their output on the existing precedents (Schaub et al., 2015). If we connect the existing research around interface reform, including the outcomes of this and future HDC QS Report Cards, we can indeed address “privacy by design” on the surface as well as the foundation of technology.

Informed Consent and Readability

The idea of policy readability is a more complex and subjective criterion, and much research has investigated potential solutions to the reader comprehension problem. The mainstream definition of human readable language stands in contrast to the high reading level and knowledge of legal vocabulary that is currently required to comprehend most privacy policies (Ermakova et al., 2014). The concept of informed consent must be shaped by the recognition that, similar to this human language literacy issue, there is a wide spectrum of technological literacy, which means that information should be comprehensible even to the newest of users (Woo, 2006).

Informed consent requires clear, comprehensible descriptions about how data will be used, as context can shift the user's level of comfort with this sort of disclosure (McCallister, Grance, & Scarfone, 2010; Lupton, 2017). In other words, the “sensitivity” of any type of collected data is much more subjective than previous practices have acknowledged, and the user should have



the agency to decide the specific contexts in which their data will be kept fully protected (Adams & Sasse, 2001; Mehlman, 2015). Use of vague language like “reasonable measures” , “material changes” and “industry security standards” are detrimental to the readability of privacy policies, and more specific terminology must be used (Acquisti et al., 2014). This would reduce a significant amount of informational asymmetry and pave the way for comprehensible modular choices (Haddadi & Brown, 2014).

Comprehensive Coverage and Consumer Education

Linguistic clarity and ease of navigation must be accompanied by appropriately comprehensive disclosure of companies’ overall practices. In particular, a company’s privacy policy needs clearly marked sections applying directly to each product or service they provide, and what information they specifically collect via those means. This is especially true for larger conglomerate companies, which frequently provide a near-useless umbrella policy which includes nothing to address the unique concerns of sensitive health information in Quantified Self products.

One of the more complex, but crucial, advances that can be made in this area is this provision of concrete and detailed descriptions of the type of information collected (Acquisti et al., 2014). There are many existing resources that outline these distinctions - a particularly strong resource is provided in *Appendix B* of McCallister, Grance, & Scarfone’s industry guide to handling Personally Identifiable Information (2010). In this document, produced for distribution by the US National Institute of Standards and Technology, they encourage companies to provide this pre-existing guide to their customers and users. This combats an often-used excuse for the lack of clear definitions- not having the resources to fully explain these issues. Even companies which may not be at the stage or scale to design detailed informative descriptions of data collection categories within their own policies can easily direct users to some of these sources to support informed decision-making.



This movement can be further expedited through collaborative cooperation within the QS industry itself. Many of the conglomerate companies which provide “user portals” on topics of privacy and security could easily allow open sharing and linking to these resources from smaller startups in the field, in a show of good faith to the overall growth of industry best practices. If this collaboration becomes normalized, it will also support alliances toward social good research.

Research for Social Good

The original aim of the Quantified Self Report Card project targeted the Human Data Commons Foundation’s central mission: to support and enable large-scale data analytics for social good research. We began this review with the hope of uncovering industry efforts to collaborate with the nonprofit and academic sectors, and provide anonymized user data to enable more effective health and well-being research. The results were disappointing: all but two companies provided no reference whatsoever to providing information without strings attached for external nonprofit projects.

[Apple’s ResearchKit project](#) and [Clue’s work to improve global female health](#) were the standout examples of what the industry could achieve by sharing its already existing, but too often siloed, reserves of human health data. Other companies provided mission and value statements that emphasized the importance of transformative effects on global health, but did not give strong indicators of “putting their money where their mouths are” outside of private sector relationships.

The possible reshaping of data sharing, and modular consent for such projects within the user interface, shifts us further from the current static focus on restricting the collection of data, to enforcing the proactive determination of post-collection usage of that data (Thierer, 2015). This would be a prime opportunity to implement new models for “opt-in” structures: at the initial stage of privacy “notice,” and within the users’ personal settings interface, they would be able to



“opt-in” to multiple avenues of such research, with flexibility around contextual factors (Vayena et al. 2013). Overall, this worldwide advance in research would benefit the companies themselves as well as users, as knowledge outcomes would be shared and contribute to rapid innovation in the health field across sectors.

Methodological Assumptions and Limitations

The review of our methodology, and potential improvements for the 2018 iteration of the Report Card, were based upon the qualitative reflections of the researchers, which can be read in full in [Appendix B](#). The 2017 Report Card will be followed by an in-depth review of the methodology’s perceived efficacy, and brainstorming for more diverse methods for next year’s edition. Some of the potential weaknesses of this year’s approach were known upon entry into the project, and some became apparent only after it was already underway.

Incorporating the large body of existing quantified research in this field would further bolster the methodological reliability of such an inquiry, though for the purposes of this initial inquiry, it can be considered a strength that the researchers approached the project without the bias of such a review. By recording their initial impressions and personal reflections upon visiting these companies’ sites before delving into a thorough literature review, they were able to more authentically represent the average user’s experience, as they will very rarely be familiar with such technical documentation.

The researchers prepared for the possibility of bias towards services they already used in their daily lives, so a category of the review form was included to indicate whether or not they had used them before. It turned out that the researchers had used very few of these companies’ offerings, but it still cannot be determined whether it impacted their qualitative perspective of those companies’ policy readability and navigability.



A qualitative methodology using only two reviewers is perceived to have very weak legitimacy within the norms of traditional formal research. The study's results are not likely to be largely replicable or generalizable. It was valuable to check the two sets of impressions against one another, and identify areas of conflicting interpretation, but this would be improved even further by more research participants.

Additionally, the experimental design focused solely on the accessibility and transparency of information within user documentation, ignoring the company's practices and standards themselves. Within the deep analysis of individual companies, it became clear upon more careful observation that some of those who scored highly on the criteria we laid out displayed red flags when it came to their actual practices. For example, many companies that did disclose how future changes would be made and whether or not users could import, export, or delete their own data got a point for informing the user of these facts, even when the information amounted to "you have no rights" or "you are solely responsible to track future changes".

Improvements for 2018 QS Report Card

The methodological advances for the 2018 edition of this Report Card will include a targeted implementation of successful crowdsourced research on the user experience and interface of these companies. The Report Card will also take into account companies that were missed in this year's review, as well as keeping an eye out for new startups. There will be a more specific breakdown of the categories and services within this sector, with research tailored to accommodate some of the differences that were noted in this review.

The Human Data Commons Foundation researchers will reach out directly to the companies that we've reviewed, directing them to the in-depth evaluation in [Appendix A](#) and seeking their feedback and correction of any misconceptions we presented. Ideally, the clarification of those "missing pieces" in navigation of public-facing documents will encourage these companies to



clarify these points on their sites as well. Reaching out directly to companies will also allow us to inquire about the more behind-the-scenes technical and security practices that are often absent from public documentation, so that we can evaluate not just what is said, but what the actual compliance with security best practices is like.

Next year's Report Card will be supplemented by a full overview and analysis of existing global regulatory frameworks for data privacy and security, from both governments and advocacy groups. By looking for the harmonious overlap between promising frameworks, we will be well-positioned to give more detailed and concrete recommendations in the 2018 review. This overview will also lead us further toward recommendations for clearer, universalized terminology in the privacy and user rights field. Ideally, we will be able to utilize these collaborative definitions to delve further into investigation about the consistency of companies' claims: by cutting through the vague surface language, we can more accurately evaluate whether there are significant divergences between their customer-facing privacy policies, and the more dense fine print of their Terms of Service.

Finally, the next year's process will stay true to one crucial element that drove this year's research-- the process will remain purposefully flexible and organic, and open to recentering when it becomes clear that our assumptions have diverged from users' self-expressed concerns.



Appendix A: Individual Company Reviews

In-Depth Reviews: Devices

Headquartered in the US, [Athos](#) produces smart performance apparel that monitors biosignals and provides analytic tools to comprehend the data. Athos received an “average” score of 2.95. The company showed particular strength with the direct and ethically aspirational statement which introduces its Privacy Policy, as well as clear language throughout. This balanced out some smaller details which show room for improvement: if the company provides a dedicated privacy officer contact, more clearly describes the ways in which future changes to policies will be indicated, and provided more specific transparency around data storage and security practices, it will be an exemplar of best practices.

Headquartered in the US, [Bellabeat](#) produces fitness-tracking wearable devices specifically targeted toward women. Bellabeat received an “above average” score of 4.4. Strong efforts were clearly made for readability, and the resulting policies are easy to understand. However, they could be further improved with more specific details about data storage and security standards. By providing a dedicated privacy officer contact, the company will be an exemplar of best practices.

Headquartered in the UK, [Blocks](#) produces modular designable smartwatches, currently in pre-order stage. Blocks received an “above average” score of 3.85. One significant strength was their thorough description of how policy changes would be indicated to the user, showing a proactive approach on their part rather than relying on the user to do the work. A point of significant concern lies in the fact that there was no separate Terms of Service document easily available, though there was reference made to a supposed link. Additionally, only one point of human contact was provided (general email address). By providing more options for direct inquiry, as well as a dedicated privacy officer contact, the company will be an exemplar of best practices.

Headquartered in the US, [Bloomlife](#) produces smart pregnancy contraction trackers. Bloomlife received an “above average” score of 4.7. The company showed exemplary results in the core grading rubric, with a well-organized and easily readable Privacy Policy, proactive standards for contacting the user about future changes, and a dedicated contact point for privacy inquiries. The only room for improvement is in the “extra credit” area, by providing more thorough and specific information on its technical security and storage practices.



Headquartered in the US, [Emotiv](#) produces wearable EEG technology for brain monitoring and cognitive assessment. Emotiv received an “above average” score of 3.85. The company displays a proactive attitude toward contacting the user about future changes, as well as utilizing easily readable language throughout much of the documentation. However, the only means of human contact is an embedded submission form. By ameliorating this, and providing a dedicated privacy officer contact, the company will be an exemplar of best practices.

Headquartered in the US, [Fitbit](#) produces fitness-tracking wearable devices. Fitbit received an “excellent” score of 5.55. This incredibly high score is perhaps to be expected given their size and early entry to the market- in some cases, the researchers have encountered “Fitbit” in use as a colloquial term for fitness trackers in general. Fitbit provides an intensely comprehensive and searchable FAQ accompanied by detailed privacy practices documentation. Since they are setting such strong standards, the one further step they could take is to give more concrete information about the technical security standards they have. Additionally, as a best practices leader, they could share the comprehensive resources they’ve produced with other up and coming players in the field, offering a link to their informational “portal” as a standard source for user education.

Headquartered in the US, [iWinks](#) produces sleep tracking wearable headbands, currently in pre-order stage. iWinks received a “below average” score of 0.25. This was the lowest score in its category group. None of the core criteria for the Report Card are met, as there is no information provided about their Privacy Policy, standards, or terms of service. The researchers readily acknowledge that this weakness is likely due to the product’s current status in pre-order stage. However, this opens an opportunity for best practices as they move toward launch- rather than needing to scrap existing policies, the company has the opportunity to embrace its “fresh start” upon original launch.

Headquartered in Sweden, [Ketonix](#) produces reusable breath ketone analyzers. Ketonix received a “below average” score of 0.95. This was the third lowest score in its category group. None of the core criteria of the Report Card were met, as no information was provided about their Privacy Policy, standards, or terms of service. Their sole strength was their sharing of three points of direct contact, something that many of the highest-ranking companies failed to do. Additionally, they received one fractional extra credit point for directly addressing the user’s ability to freely import and export data, a topic that had to be extrapolated indirectly from the policies of many other companies. As Ketonix is based in Europe, they face the opportunity to embrace the strong data privacy practices that have traditionally been enforced in the European context. It must begin with providing at least a basic Privacy Policy, however.



Headquartered in Canada, [Mio](#) produces fitness-tracking wearable devices. Mio received an “excellent” score of 5.15. The researchers both felt there was significant effort displayed in providing relevant information in clear and easily readable language, and their legal documentation included non-required information educating about digital security practices and user-side options to ensure data privacy. Mio displayed some of the most impressive practices in encouraging users’ agency, including a section in the Privacy Policy titled “Your Choices” which breaks down all possible privacy settings, and what they involve. Additionally, the company gives clear warning that the users’ data will potentially be transferred across US soil or in other international contexts. Mio could be a beacon of best practices by providing a dedicated privacy contact, clarifying the statement that the company “may” inform users proactively about policy changes, and more clearly defining why they seem to have two privacy policies available.

Headquartered in the US, [Misfit](#) produces fitness-tracking wearable devices. Misfit received an “excellent” score of 5.15. The company meets all core standards of the Report Card and provides multiple points of human contact. They display best practices by delineating different privacy policies for the website and the actual application, something that is often unclear in others’ privacy policies. They transparently address the fact that there are limitations to full deletion of personal information, and that de-identified data will not be deleted. They stand out by providing a statement self-identifying their compliance with Privacy Shield standards, which are still being transitioned into full implementation in the public sphere. There is room for improvement insofar as their incredibly long documents could be made easier to navigate, perhaps with a hyperlinked and fixed table of contents.

Headquartered in Canada, [Muse](#) produces cognitive feedback-based meditation headbands. Muse received an “excellent” score of 5.35. This was the third highest score in its category group. The company meets all core standards of the Report Card and provides multiple points of direct contact. The Privacy Policy was very clear, thorough, and direct, and they were the first company noted to differentiate the data they collect into three distinct categories, ranging from most to least sensitive. They explicitly state that users own their own data. They also provide a very clear description of how the data is locally hosted on the user’s device, and what that means for its security. They have a somewhat vague statement that a “reasonable effort” will be made to notify users about future Privacy Policy changes, so it would be of benefit for them to define what “reasonable” means. Additionally, they could benefit from more specifics on their technical security standards, not necessarily because they are comparatively weak in this area, but because it would bring their “extra credit” ranking into line with their other above average



documentation. Finally, they should provide a dedicated link for each document, as it can be slightly confusing to navigate the current layout.

Headquartered in the US, [Nervana](#) provides devices which synchronize nerve stimulation with sound, allowing the user to experience relaxation through technology. Nervana received an “excellent” score of 5. The company meets all core standards of and provides multiple points of direct contact. Though there is a lack of clarity in some of the documentation, it seems to be an issue with the resources they have in their current state of development, because what they have seems to be user-oriented. However, the policy’s content is mostly generic, rather than focusing specifically on customer’s potential experience with their devices. It was concerning to the researchers that there is no visible Terms of Service link, although the traditional content of such a policy was addressed in a generic “Legal” link. Additionally, they place the responsibility to track changes in the Privacy Policy solely on the user. Finally, Nervana was one of the few companies to receive no extra credit points, so future improvements could focus on providing more information about their technical standards and storage of data. Interestingly, the company appears to be an outlier example in this Report Card, as it technically succeeded at most of the rating criteria but a more qualitative and subjective inquiry showed that its documentation leaves much to be desired, thus apparently “just barely passing” in each category of review.

Headquartered in Finland, [Polar](#) produces fitness-tracking wearable devices. Polar received an “average” score of 2.7. The researchers did not get the impression that attempts were made to provide a user-friendly, readable Privacy Policy. Additionally, the Privacy Policy and Terms of Service were subsumed under a sole “Legal” footer link, and there was no dedicated privacy contact provided. However, there were a few unique strengths in their documentation. They provide thorough documentation of what categories of data are specifically contained in their files, to a degree that was not seen in any other reviewed companies. Additionally, extensive documentation of their ethical stance and practices regarding environmental responsibility show the possible staging ground for them to provide the same best practices regarding user rights and privacy.

Headquartered in Poland, [Sidly](#) produces multi-tracking wearable devices for home healthcare environments. Sidly received an “average” score of 3.2. The researchers did not get the impression that attempts were made for a user-friendly or easily readable Privacy Policy, and no direct privacy contact was provided. However, they provided multiple points of direct contact for user inquiries, and the website’s fixed footer clearly delineated each informational link that was required for navigation of their policies. However, there appeared to be a language barrier within their policies - the existing English language documents needed more effective translation, and



would benefit from being shortened or simplified. More concerning is the fact that, within the English language website, the links to the Privacy Policy and other detailed information led to the Polish language documents. However, the strengths demonstrated in Sidly's thorough and navigable layout suggest that, by solving these translation issues, they could emerge as a paragon of best practices.

Headquartered in Germany, [Sigma Sport](#) produces fitness-tracking wearable devices and heart monitors. Sigma Sport received an "excellent" score of 5.4. This was the second highest score in its category group. The researchers found it quite easy to find their dedicated data statement, which was clear and well structured. They also provided multiple points of direct contact for user inquiries. Perhaps their strongest practice was their vocal compliance with the German Federal Data Protection Act, one of the strictest sets of guidelines around security measures and confidentiality. Though of course this is due to their operation on German soil, it allows them to serve as an example of a best practice company on a global level. However, this could be further improved by translating the additional documentation provided into English alongside the existing German.

Headquartered in the US, [Soleus](#) produces fitness-tracking wearable devices, targeted specifically toward runners. Soleus received an "average" score of 2.95. The researchers did not get the impression that attempts were made to provide a user-friendly and readable Privacy Policy, as it seemed to be an enactment of a simple boilerplate policy template, and the Terms of Service could not be easily found. The company would also benefit from providing a dedicated privacy contact. However, they demonstrated the strengths of providing multiple points of direct contact, and a transparent description of the caveats of their storage of data on US soil.

Headquartered in Finland, [Suunto](#) produces fitness-tracking wearable devices. Suunto received an "above average" score of 4.9. The company met all core standards for the Report Card requirements, as well as providing one of the more specific descriptions of their security standards which was found among the companies reviewed. They also demonstrate transparency when describing the limitations on fully deleting data from their servers, and describe the best actions to ensure data security and privacy on the user side. Their posted statements of conformity to industry standards and dedication to protecting the environment can easily be translated to taking a future stand on user privacy rights. Additionally, there is room for improvement in the ease of readability of their documentation, and simplifying the site's navigation and user interface.



Headquartered in China, [Tic](#) produces fitness-tracking smartwatch devices. Tic received a “below average” score of 0.7. This was the second lowest score in its category group. The researchers did not get the impression that attempts were made to provide a user-friendly and easily readable Privacy Policy. Additionally, there was significant concern insofar as the Privacy Policy was referenced within the Terms of Service but doesn’t seem to actually exist. In line with this, there was also no dedicated privacy contact listed. The documentation was hard to navigate, lacked many crucial elements, and what was available was difficult to read. However, The researchers acknowledge that this may well be due to a language barrier, as much of the documentation seems to be translated. Greater attention to detail and a more careful translation would significantly improve this company’s rating.

Headquartered in the US, [Wahoo](#) produces fitness-tracking wearable devices and applications. Wahoo received an “above average” score of 4.2. The Privacy Policy they provided seemed direct and easy to read, and they provided multiple points of direct contact. However, there was no dedicated privacy contact listed, and the researchers were not able to find the Terms of Service agreement. Additionally, they indicate that the user is responsible to track future changes in the Privacy Policy. Though they provide information about their security practices, they are somewhat generic, and their documentation would be improved by better defining concepts such as “appropriate” or “reasonable” steps to ensure data privacy and security.

Headquartered in France, [Withings](#) produces fitness-tracking wearables and other devices. Withings received an “excellent” score of 5.25. The company meets all core criteria for the Report Card review, including a dedicated privacy officer. Their policy documentation is remarkably strong in its emphasis on users’ rights to control, import, export, and delete their own data, and it goes into thorough detail about what data they collect and why. They take the proactive step of informing users directly of future changes to the Privacy Policy as well as provide a standalone Data Protection Policy to accompany the Privacy Policy. They disclose significant detail about their security practices and storage of data. One area for improvement is the company’s provision of direct contact points, as they only provide the physical addresses of their various geographical locations. Instead of an email address or phone number, the user must click through a labyrinthine FAQ section in order to even reach a “contact us” section. This is particularly problematic for users who might have general inquiries, in which the options of the FAQ serve little use.

In-Depth Reviews: User Platforms



Headquartered in the US, [Beeminder](#) is an online and mobile platform for escrow commitment contracts to build user-designated habits. Beeminder received an “above average” score of 4.6. Beeminder does not provide a dedicated privacy contact, and the company holds the user responsible to track future changes to its policies. However, these policies displayed consistently clear and easily readable language, presenting their mission, values, and practices without excessive use of legal terminology. Beeminder is unique in its utilization of tongue-in-cheek humor interspersed with more formal legalese in its documentation. The company provides four points of direct contact for general inquiries, and emphasizes users’ complete control over their data and its export or deletion. Interestingly, they disclose that their documentation is based upon [Wordpress’s open-sourced terms of service](#), which could serve as a forward-facing example of informational resource sharing between companies which benefits the end user. Ultimately, Beeminder’s documentation gave the researchers the impression that they put thought and effort into communicating their philosophy about data privacy, rather than just their operational practices.

[BioBeats](#) produces well-being and coaching products for health and productivity. BioBeats received an “above average” score of 3.85. This was the only company that did not clearly indicate where its headquarters are located. Though BioBeats provides a dedicated privacy contact, they hold the user responsible to check for future policy changes, and an aggressive disclaimer that if the user disagrees with future changes they must cease using the application immediately. Their documents are well laid out, with an easy to navigate structure, but the language is dense and not easily readable. They provide clarity around the services themselves, but very little transparency regarding their privacy practices. However, they do provide informative links to the Privacy Policies of third parties that they may share user data with, such as Google Analytics. Though they do not list specific security standards used in their operations, they provide a comprehensive list of the tools and services used to process data in practice within their Privacy Policy. They do provide comprehensive information on the users’ rights to request exportation or deletion of their personal data.

Headquartered in Germany, [Clue](#) is a mobile app for reproductive health tracking. Clue received an “excellent” score of 5.3. The company provides a dedicated privacy contact (although it is not prominently displayed within the policy), two points of direct contact for general inquiries, and indicates that it will inform users of any future “material” changes to the Privacy Policy (without specifically defining this term). The Privacy Policy is laid out in an easily navigable format.. The company displays a consistent commitment to explaining both the ethical side of their data collection policies, and the technical protocols used to protect user data. One crucial strength that sets Clue apart from other companies is the option to utilize their service without



ever setting up an account, or submitting personal information - the application can run locally on a device without requiring data transmission to and from the cloud. Additionally, the company indicates that for those who do set up an account, their personally identifiable contact information is stored separately from their health tracking and service data. The company explains the cryptographic protocols by which passwords are protected, the secure 'https' protocol used for access, and steps users can take to ensure data security on their side. Combined with the CEO's values-centric public blog post on their work for global female health, Clue is a paragon of potentially best practices. The primary recommendations for improvement would be a more prominent listing for privacy inquiries, including a Privacy Officer rather than merely a general "data/privacy" email address, and clarifying their definition of "material" changes to the Privacy Policy .

Headquartered in Switzerland, [Dacadoo](#) is a platform providing a health score measurement for longitudinal analysis. Dacadoo received an "above average" score of 4.6. The company does not provide a dedicated privacy contact, and the user is held responsible to check for any future changes to its policies. Though there is no overall summary provided within the Privacy Policy, it is written in relatively easy-to-read language, with an in-depth "Frequently Asked Questions" section and coverage of privacy and security information. The company also provides four points of direct contact for general inquiries. Though Dacadoo describes that it follows the "high degree of data protection regulation" in Switzerland, it does not give detailed information on what these standards involve, leaving it to the reader to find out more. However, this declaration allows them to indicate that they can only guarantee these standards are upheld until the point the services are accessed from other countries, at which point the data will cross borders and be subject to potential compromise. This sort of disclosure is a crucial aspect to transparency and informed consent within user-facing documentation.

Headquartered in the US, [Fitness Syncer](#) is a unified dashboard for users to import and organize health and fitness data. Fitness Syncer received an "above average" score of 4.2. The company does not provide a dedicated privacy contact, and they indicate that they will notify the user with "material notices" of policy changes, without clearly indicating what would be considered a material change. They provide transparent user-facing information about all of the partnerships they conduct, including the form of data transfer (whether uni-directional or bi-directional) that they engage in with these platforms. Their overall Privacy Policy consists of easy-to read language, and they provide four points of direct contact for general inquiry. They acknowledge the user's right to close their account, but indicate that they will retain certain information even after accounts have been shut down, and they provide a vague description of their utilization of "standard" security practices.



Headquartered in the US, [HeartMath](#) produces mobile and online applications for self-directed anxiety reduction practice. HeartMath received an “above average” score of 3.95. The company provides a dedicated privacy contact (listed as the VP of Finance and Operations, but indicated to be available specifically for such inquiries), and holds the user solely responsible to check documentation for future changes. There was no visible Terms of Service agreement, and the Privacy Policy available was very dense and difficult to read language. They indicate that customers can request a deletion of their information “in most cases” by contacting Customer Service, without explaining the exceptions. Additionally, they have detailed information about security practices, but only seemingly around the collection and transmission of credit card information. They provide three direct points of contact for general inquiry, but the researchers got the impression that informing potential users of their standards and practices was not a particularly high priority.

Headquartered in the US, [LibreView](#) is a cloud-based diabetes management system that provides reports from many popular glucose monitoring devices. LibreView received an “above average” score of 3.9. The company provides a dedicated privacy officer with multiple avenues of contact, but holds the user responsible to check for future policy changes. Though relatively well-organized, the policies are of dense legal language, and it is difficult to understand which policies apply to which LibreView services, as there seem to be multiple portals for patients, professionals, and use of the site itself. However, they provide a thorough and exemplary Frequently Asked Questions section which addresses many key issues around data privacy and technical security. They provide two points of direct contact for general inquiries, which are prominently displayed in multiple locations on the website.

Headquartered in the US, [My Fitness Pal](#) is an online and mobile calorie counting and nutrition platform with fitness service integrations. My Fitness Pal received an “above average” score of 4.45. The company provides a dedicated privacy contact, but indicates that they “may notify” the user of future policy changes, without indicating why or why not. The documentation they provide is thorough but labyrinthine and dense with legal terminology. They use surprisingly specific case scenarios in their Privacy Policy, linking multiple “read more” human language descriptions to better illustrate the more general terms. It is crucial to note that, as an apparent acquired property of Under Armour, the My Fitness Pal documentation redirects to Under Armour’s site and policies: this is particularly important as we review the latter separately under the Conglomerate category. This redirect makes it difficult to find specific information that relates to the My Fitness Pal platform itself. Information security practices are addressed, but the description is vague, and they disclose that European citizens have the right to request permanent deletion of their data without explaining the rights of non-European citizens. Overall, these policies provide relatively thorough and transparent documentation, though it is



somewhat burdensome to navigate and difficult to uncover the information that may be specific to My Fitness Pal.

Headquartered in the US, [Nutra Hacker](#) is a platform for users to upload DNA information for specialized analysis and reports. Nutra Hacker received an “average” score of 2.65. The company does not provide a dedicated privacy contact, nor do they disclose any information about how future changes to their policies will be indicated to users. Though the company can be commended for using easily readable language in its Privacy Policy, it lacks significant detail. It is the shortest Privacy Policy in our review, comprised of a mere six sentences. Even in its brevity, the documentation shares a small amount of specific information about Nutra Hacker’s technical security practices, disclosing that traces of potentially identifiable information are “typically” deleted each month to reduce the risk of matching users’ health and browsing information. The researchers were unsure whether there was a means to access further details about the company, especially as the open discussion forums were inactive, and there was only one point of direct contact given for general inquiries.

Headquartered in Israel, [Nutrino](#) is a nutrition insights platform providing personal health recommendations. Nutrino received an “average” score of 3. The company does not provide a dedicated privacy contact, and places responsibility upon the user to check for future changes to its policies. Their Privacy Policy, though not extremely informative, is written in direct and easy to parse language. They also provide four points of direct contact for general inquiries, and upon receiving a written request to delete user information, they will do so fully so that it cannot be restored; however, they include a standard clause about retaining certain information when necessary, and they do not specify what type of information this would be.

Headquartered in the US, [Pacifica](#) is a stress and anxiety self-tracking and management mobile app. Pacifica received an “average” score of 2.85. The company does not provide a dedicated privacy contact, but they will notify users in case of “material changes” to their policies, without defining what would constitute such a material change. The Privacy Policy is written in dense legal language, and would benefit from one easily readable summary of its overall content. There is a scant, but existing, description of their data security and server storage practices, and one point of direct contact for general inquiries. They do reference the option for users to delete their account and data, though they include a standard clause explaining that it may not be possible to remove all data due to technical, legal, or contractual restraints. Additionally, they proactively address the issue of data tracking on their “About Us” page, in addition to the Privacy Policy.



Headquartered in Australia, [Predict BGL](#) is a platform providing predictive analytics for diabetes management. Predict BGL received an “above average” score of 4. The company provides a dedicated privacy contact, but fails to describe how future changes to its policies will be indicated to the user. The policies are long and somewhat difficult to read, but structurally well-organized and impressively thorough in terms of technical disclosures. They provide four points of direct contact for general inquiry, and overall their review of not only their technical practices, but the implications and potential risks of cloud data storage and cross-platform sharing exhibit best practices in pursuit of informed user consent. Additionally, they discuss HIPAA compliance, which was a rare acknowledgement among the companies we reviewed.

Headquartered in Great Britain, [SAMapp](#) is an anxiety self-tracking and management mobile app. SAMapp received an “above average” score of 3.65. The organization provides a dedicated privacy contact, naming a specific individual in the role of Data Controller, but fails to describe how future changes to its policies will be indicated to the user. SAMapp provides a notably succinct Privacy Policy, composed of a single page and containing the most easily readable Terms of Service Agreement we encountered in this review. They provide only one point of direct contact for general inquiries, but their technical disclosures and acknowledgement of compliance with privacy and security regulations is exemplary. It seems likely that these strengths can be traced to the application’s origin in a University research context; as the organization seemed to pursue this project in a non-profit context, and bound by the somewhat rigid ethical standards of academia, the creators likely face much less pressure to focus on long-term liability reduction and maintenance of proprietary techniques, being driven instead by motivation to educate end users. They disclose the specific type of database upon which information is stored, as well as their Internet Service Provider, and they make direct reference to compliance with the 1998 Data Protection Act.

Headquartered in the US, [Training Peaks](#) is an eco-system of web, mobile and desktop products for endurance event training. Training Peaks received an “average” score of 3.3. The company does not provide a dedicated privacy contact, and though their description of future change notification is quite thorough, it places responsibility for tracking policy changes upon the end user. However, their choice of language seems to encourage the user to contact them with any questions about these changes, and it is the only policy we reviewed which designates a thirty day “grace period” between posted notice of changes to the Privacy Policy, and the enactment of those changes. Overall, the policies are written in rather dense legal language, although a short summarizing section of the Privacy Policy is written in more accessible language. They provide a Frequently Asked Questions section that, while well written, is poorly organized. They provide two points of direct contact for general inquiries, and provide a thorough statement of their dedication to users’ privacy rights, although its description of best practices is somewhat



generalized. It is encouraging that they remind the user, at the very end, that even with their strongest efforts, “perfect security” does not exist, and that there are some cases in which these efforts may fail.

Headquartered in the Netherlands, [X2AI](#) produces Tess, a psychological AI that administers personalized health-related reminders. X2AI received an “excellent” score of 5.15. The company provides a dedicated privacy contact and names a specific Security Officer, though it holds the user responsible to check for future changes to the policies. However, X2AI’s Privacy Policy explains the company’s practices and values in user friendly language, with a particularly detailed and easy to understand section on security. This security policy was among the most thorough and readable across the entire body of reviewed companies. In particular, their thorough transparency regarding rigorous technical standards, including their support of users’ rights to access their services with Tor and their disclosure of specific high-level encryption practices, went above and beyond other policies we encountered. The main room for improvement lies with their two points of direct contact for general inquiry, which the researchers could only find by looking in their white paper– these should be displayed prominently within the site itself.



In-Depth Reviews: Middleware Analytics

Headquartered in the US, [Advanced Brain Monitoring](#) provides EEG and sleep monitoring devices and software for professional use. Advanced Brain Monitoring received an “average” score of 3.15. The company does not provide a dedicated privacy contact, nor describe how future changes to their policies will be indicated, but they do provide three points of direct contact for general inquiries. Their Privacy Policy is straightforward, written in user-focused language, and contextualizes its content in terms of concrete use of their service, rather than generalized operations of the website. They do not provide explicit technical details about security or storage practices in the basic user-facing policies, but with extra work, it is possible to track down the comprehensive documentation for the specific Sleep Profiler service, which includes detailed technical information about compliance to industry standards.

Headquartered in the US, [Affectiva](#) provides emotion recognition AI technology. Affectiva received an “above average” score of 3.95. Affectiva provides a dedicated contact for privacy inquiries, as well as three points of direct contact for general inquiries. They display a rare option to opt-in to direct email notifications about future policy changes, which removes responsibility from the user to check the website regularly, but also does not force them to receive such emails. Much of the documentation is in complex “legalese” language which may be difficult for end users to navigate. However, this documentation provides impressive descriptions of their dedication to user privacy, disclosing that they retain the minimum amount of anonymous technical data that is necessary for operations, and that all specific data processing is done locally on devices. They encourage developers to follow similar best practices, which is crucial for the Middleware category.

Headquartered in Israel, [BeyondVerbal](#) provides vocal emotional analytics. BeyondVerbal received an “average” score of 3.2. BeyondVerbal does not provide a dedicated privacy contact, and the company leaves the user responsible to check the policies for any future changes. The Privacy Policy is brief, vague, and not particularly descriptive. It is not clear if they collect data themselves or just provide the algorithms/tools for others to do so. The policy is displayed as one large block of text, rather than being segmented by topic headings, and it appears to apply solely to the website rather than any of the company’s services. The one strength demonstrated by BeyondVerbal was the company’s provision of four points of direct contact for general inquiries.



Headquartered in the US, [Eyeris](#) provides deep learning-based emotion recognition software that reads facial micro-expressions. Eyeris received an “excellent” score of 5.2. The company provides a dedicated privacy contact, as well as providing four points of direct contact for general inquiries, but holds the user responsible to check for future policy changes. Eyeris makes multiple explicit statements regarding their commitment to users’ privacy. These statements focus in particular on the fact that, aside from specific circumstances in which customers reach out directly with service issues, the company avoids collecting any personally identifiable information in order to provide its services, and explicitly prevent their corporate customers from sharing such information with them on the backend. The sole red flag in this documentation is a relatively simple fix: the Privacy Policy hyperlink within their Terms of Service document is a broken link which needs to be redirected.

Headquartered in Canada, [Fatigue Science](#) utilizes biomathematical science to provide insight into sleep and fatigue from wearable device data. Fatigue Science received an “average” score of 2.95. The company does not provide a dedicated privacy contact, nor do they explain how future changes to their policies will be indicated, as they only comment that they retain the right to amend the terms at any point. Their documentation is relatively straightforward, though this may be in part because they do not address services in detail, instead focusing on use of the website. The policies are laid out clearly with section headers that make it easy to navigate, and they provide three points of direct contact for general inquiry. However, there are a few points of strong concern, which include their statement that they may charge a fee in order to fulfill data deletion requests, and their warning that they are under “no obligation to protect or secure any information” which the user provides when using their website.

Headquartered in the US, [Human API](#) provides a real time data health network for the aggregation and analysis of professionals’ health information. Human API received an “average” score of 2.9. The company does not provide a dedicated privacy contact, and leaves the user responsible to check for future changes in its policies. Unfortunately, the documentation is in dense language emphasizing limited legal liabilities rather than user rights. They utilize clear language when discussing their stated values, but provide little clarity regarding their actual practices. On the positive side, they provide two means of direct contact for general inquiries, and provide separate Terms of Service documentation for company relationships and end users. The company describes the customer’s right to request account deletion and the deauthorization of future data collection, though there is not a direct confirmation that such requests will be reliably fulfilled.

Headquartered in Denmark, [iMotions](#) provides eye tracking, facial expression analysis, galvanic skin response, and other analytic technologies. iMotions received an “average” score of 3.35.



The company does not provide a dedicated privacy contact, and only describes how future changes to the Terms of Service will be indicated, with no commentary on this in the Privacy Policy. They provide three points of direct contact for general inquiries, and their overall documentation is highly informative and educational. In particular, they describe their security practices with more clarity and transparency than most other companies surveyed. However, their “Help Center” documentation is only readily accessible to existing customers. Nonetheless, they do provide an explanation of how to request an exception to this access restriction by contacting them directly, although such access is granted “on a case by case basis.” Without knowing the content of this Help Center, it is difficult to comment upon its usefulness.

Headquartered in the US, [Qusp](#) provides a platform for biosignal interpretation for integration into existing mobile and desktop applications through a cloud API. Qusp received an “above average” score of 3.75. The company does not provide a dedicated privacy contact, and holds the user responsible to check for future changes to its policies. The Terms of Service documentation was more reader friendly than most, and the Privacy Policy utilized easy to understand language and maintained brevity as needed. The company provided three points of direct contact for general inquiries, and the fact that their services are currently in a beta stage makes their excellent documentation impressive.

Headquartered in Canada, [Sensaura](#) provides emotion detection and analysis technology based on data from wearables. Sensaura received a “below average” score of 1.15. The company does not provide a dedicated privacy contact, nor do they describe how future changes to their policies will be indicated. This is perhaps in part because these documents do not exist: there is a Terms of Service agreement for developers, but no user facing documentation. However, the documentation they do provide for developers includes impressive expected standards for how those developers handle user-facing data practices, requiring them to ensure they will “use commercially reasonable efforts to protect” personally identifiable information that they collect from end users, as well as refraining from exposing that information to third parties without users’ explicit consent. This highlights a flaw in our projects’ methodology-- by the scoring criteria we defined, Sensaura scores very low, but their practices seem aspirational in terms of user rights-- they just are not operating in a direct user-facing capacity. Their strong dedication to ensuring their developers engage in best practices could be highlighted in easier to access public-facing documentation to make this information more clear.

Headquartered in the US, [Validic](#) is a storage and analytics platform for aggregate user health data. Validic received an “excellent” score of 5.6. The company provides a dedicated Privacy Officer with an email and physical address, but the user is held responsible to track future changes to their documentation. The policies they provide are readable and well organized,



though somewhat long and reliant upon legal and technical terminology. They impressively provide a separate Data Security Policy with some of the most thorough documentation that we uncovered in this project, including a detailed breakdown of each of the types of data they collect and what they use it for. They provide four points of direct contact for general inquiries, and self-proclaim their compliance with both the EU-US Privacy Shield framework and the US-Swiss Safe Harbor framework. Overall, the researchers got the impression that Validic strongly prioritized informed user consent about security and privacy practices.



In-Depth Reviews: Conglomerates

Headquartered in the US, [Apple's iOS Health](#) is a fitness-tracker service platform which consolidates data across Apple devices and integrated services. iOS Health received an "above average" score of 4.3. The company does not provide a dedicated privacy contact, and responsibility is left to the user to check the site for policy changes. The company provides thorough, in depth plain language coverage of privacy issues across the board, though the content is somewhat dense at times. One point of concern is that the provided Privacy Policy is an overall document for Apple as a company, without a specifically tailored Policy for the Health service. Additionally, the information provided about their security standards is vague, and could be clarified. Despite these limitations, Apple displays aspirational vision in its overall summary of privacy practices, directly addressing the user's agency in simple language. The company indicates that it prioritizes flexible and modular user choice regarding which information is stored in the Health platform, which applications are provided access, and that the user's private passcode will control the encryption of all data. Apple also encourages users to read the individual Privacy Policies of all external companies before granting them confirmative access to this data. Finally, Apple was the sole company in our review which provided a prominent, dedicated option for contributing anonymized data to research for social good, an initiative that has earned them accolades within the larger technology field.

Headquartered in Switzerland, the [Garmin Vivo Smart 3](#) is a smartwatch activity-tracking device. Garmin received an "above average" score of 3.8. Garmin provides a dedicated privacy contact, and takes the unique approach of providing a [security and vulnerabilities submission form on a standalone page](#). Their greatest strength is their full disclosure of the possible risks of data vulnerability or surveillance in their larger global operations. In particular, they describe in detail that when data is shared with companies outside of Switzerland and the European Economic Area it is done so in accordance with European Commission Model Contractual Clauses. Additionally, they describe the technical processes by which data is synced from users' local devices to Garmin servers. Despite thorough documentation, they would benefit by making it easier to read, and they require the user to track future changes to the Privacy Policy. Their greatest weakness is the absence of information on direct contact for general inquiries, instead proving a labyrinthine set of Frequently Asked Questions for the user to navigate.

Headquartered in the US, [Google Fit](#) is an application and platform for tracking fitness, health, and activity movement. Google received an "above average" score of 3.8. Given their ubiquity in the technology market, it was surprising to see that Google lacked a dedicated privacy contact,



as well as providing users with no means for direct contact in general. Despite this, they provide thorough, direct, and easily readable documentation about their larger privacy and security practices. This includes transparent disclosure of their operations around tracking and targeted advertising. However, these resources operate as umbrella policies, with no apparent specifics addressing the Fit platform itself. Their privacy portal resources can be seen as a stellar example for the industry's movement toward informed user consent, especially if they choose to make these available for wider distribution and use.

Headquartered in the US, the [Microsoft Band](#) is a health-tracking smartwatch. Microsoft received an "average" score of 3.45. The company lacks a dedicated privacy contact, and would benefit from clarifying their processes of notifying users of future Policy changes, as they currently state that users "may be" notified. They provide thorough documentation on the company's overall privacy principles within their "[Trust Center](#)," but they provide very little specific information about how this extends to health data and wearables. They do disclose some information about where and how data is stored, though not very specific, and they disclose that users can delete some data, but do not clarify which type of data cannot be deleted. They would benefit by more clearly organizing the navigation of their privacy and security documentation, given its vast scope. Additionally, their "[Contact Us](#)" page involves an unusual virtual assistant interface, making it difficult to access the single means of direct contact provided beyond that interface.

Headquartered in the US, the [Motorola 360 Sport](#) is a fitness-tracking smartwatch. Motorola received an "above average" score of 4.85. Motorola has a dedicated privacy contact, but holds the user responsible for tracking changes to its Privacy Policy. Some of the company's documentation was perceived as utilizing easily readable language by the researchers, but this documentation was long and not specific to any of the products or services provided. Additionally, it was difficult to determine what part of this documentation was the Terms of Service, as the divisions between different elements were not clearly demarcated. Perhaps the weakest point of Motorola's practices was the elusive display of a direct means of contact: it was necessary to click through approximately eight layers of a "troubleshooting process" interface before accessing the option for their sole email address for general inquiries. As far as strengths, Motorola utilized a more specific disclosure of technical security standards than most other companies, with transparent emphasis on the fact that their encryption protocols cannot entirely guarantee the protection of user data due to its travel through various external channels. It's perhaps not the most user-friendly approach to end with "you use our site at your own risk," but this protective disclaimer is at least honest about the imperfect methods of online data security.



Headquartered in South Korea and the US, the [Samsung Gear Fit2](#) is an activity-tracking smartwatch. Samsung received an “above average” score of 3.75. Samsung did not provide a dedicated point of contact for privacy concerns, with a link in their Privacy Policy redirecting to the generic “general inquiries” page, and the company places responsibility for tracking policy updates on the user. Their documentation was presented with a confusing navigation structure, but the content was thorough and straightforward. Though they provide clearly sectioned and sub-titled sections, the policy itself utilizes complex language and is long, as well as covering the general operations of the company and its website without directly addressing any of the products or services. They address technical security procedures in a generalized manner, claiming to use “reasonable measures” for data protection without clarifying what those measures are. However, Samsung provided the most points of direct contact out of any of the Conglomerate companies on their general inquiries page.

Headquartered in Canada, [Telus HealthSpace](#) is a platform for online health information storage. Telus received an “excellent” score of 5.1. Telus provides a dedicated privacy contact, as well as two points of direct contact for general inquiries. They also indicate that they will contact the user directly with any “material changes” to their policies. They address their security and encryption practices, although the description is still somewhat generalized, and they do not provide specific details about whether they merely store or actively synchronize user data with the integrated external sources. However, they do disclose explicit details about the practices used in on-site server storage. The company provides some of the most thorough and direct documentation of all companies reviewed, but it would benefit from providing summaries for each section and built-in hyperlink navigation, given its extensive length. They explain not only the type of data they collect, but also break down the definition of what each type of data includes; however, the data collection addressed seems to solely refer to use of their website, with no specific information about their collection of health data within the platform. They provide one of the most explicit explanations of the process of data deletion, allowing permanent deletion of information within ninety days after a user request. Overall, Telus displays an impressive attempt at best practices in regards to transparency, security, and informed user consent.

Headquartered in the US, the [Timex Move x20](#) is a fitness-tracking smartwatch. Timex received an “average” score of 2.5. Timex does not provide a dedicated point of contact for privacy inquiries, and holds the user solely responsible to track policy changes. The researchers did not feel there was an attempt made to produce a clear or easily readable Privacy Policy, and the documentation they provide only applies to the website, rather than specific products or services. After a labyrinthine search for information about the data practices and standards for their wearable products, the researchers could not find such disclosures, even within the



provided user guides for those products. While they provide two points of direct contact for general inquiries, the overall documentation seemed more focused on limiting company liability than on informing users of their rights. They provided transparent warnings, about their inability to guarantee data would not be disclosed under specific circumstances such as security breaches and government requests. Since this is a technical reality for all data collection, it was considered to be a sole strength in their documentation practices, especially as its language encouraged users to seek further information about these practices in the general technology field.

Headquartered in the Netherlands, the [TomTom Touch](#) is a wearable fitness-tracking device. TomTom received an “above average” score of 4.2. TomTom provides a dedicated contact for privacy inquiries, listing an entire “Privacy Division” within the company, although they don’t provide points of direct contact for general inquiries. They indicate how future policy changes will be indicated within the Terms of Service, rather than the Privacy Policy, and responsibility is put upon the user to track these potential changes. The Privacy Policy is in easily readable language, and in many cases shows a proactive focus on the user’s rights, rather than merely limiting the company’s liability. In particular, the policy discloses that the user retains the right to withdraw permission for use of their data at any time, as well as the right to import and export data to and from other sites. The deletion of data will be guaranteed within thirty days of the user’s request, and cannot be restored, indicating that best practices are made to completely eradicate such data internally. TomTom indicates that in general, all personal data will be processed in the European Economic Area, and when it is necessary to transfer data outside of this region, they will ensure “adequate safeguards” to protect the user’s privacy rights. Additionally, TomTom indicates that any further use or processing of user information that is not covered in these policies, other than those necessary for essential data storage purposes, will not be carried out without direct notice to and informed consent from the user. Finally, they claim to enforce similar standards of user privacy protection in their collaboration with external companies. All of these practices are impressive examples for the larger industry, although the documentation could benefit from including summaries and easier navigation.

Headquartered in the US, the [Under Armour Gemini 2 Record](#) is a shoe containing an activity-tracking device. Under Armour received an “excellent” score of 5.1. Under Armour provides a dedicated point of contact for privacy inquiries, as well as two points of direct contact for general inquiries. They state that they “may notify” the user of changes to their policies, without specifying determining factors. The content of their policies is very dense overall, but the structure utilizes a well-organized sidebar utilizing an outline table of contents. There are also highlighted summaries in everyday language, although these seem at times to contradict the long-form legalese descriptions. The Privacy Policy indicates that the user will be



provided with options to opt-in or opt-out of data collection upon setting up an account, but it also indicates that they own all user-generated content and will retain it even if the user deletes their account. Their documentation is very strong in terms of country-specific disclosures, and they provide a somewhat vague but overall educational discussion of data security, informing the user on how to keep their own data safe.



Appendix B: Researcher Process Reflections

Reflections: Rochelle Fairfield

Concrete Research Process

Researching the Privacy Policies and Terms of Service for platforms, services and devices engaged my analytical and cognitive capacities in a way that produced a kind of altered state. I'd feel my mind come alive in the stretch between scanning the website, taking in the Gestalt, then dropping into the zone as I narrowed in on the particulars of the Privacy and Terms of Service text. It was rewarding of itself to get into this zone, and feel a kind of intimacy with the legal principles that underlie the domain of Quantified Self services, platforms and devices.

Reading the Terms of Service was like hearing a dialect of English (my first language) that I recognized, yet had to think about to get the meaning. Most Privacy Policies used language that contrasted significantly with that used in the Terms of Service. The Privacy Policies were generally more conversational in tone, sometimes even an impassioned proclamation of being for the user/reader. Meanwhile the Terms of Service were generally passionless in tone, and a one-way dictation of who owns what. At a certain point each day, after reading and searching about 3-5 websites and docs, I could feel my concentration start to waver, and used this prompt to pause and refresh body and mind. Reflecting on this now, I smile thinking that there's probably a device that could track my brainwaves during research and dashboard my brain-waves and corresponding concentration: note made to research the researcher for next years' Report Card!

What I Realized as I Went

Something I started to realize during the research is that using these devices creates data, direct and passive, and companies harvest this data for numerous reasons. One is to provide their



service, fair enough. But reading some 50-ish docs it became apparent that there is a practice of claiming the right to also use data for purposes other than providing the product or service.

Future use of the data, including having it be part of a company's assets in the event of sale or acquisition, points to myriad implications of collecting this directly relevant, as well as non-essential data. Relative to the promise of the Privacy Policy to respect a person's autonomy, many of the companies have a competing commitment of getting as much data as possible from people so as to be able to capture value from it in known, as well as not yet imagined ways. This creates a tension between providing a product or service for the client, and using the client to harvest information that can then be used in a different marketplace, that of options trading, data brokering, and the like. This incentivizes devoted use of, addiction even, to the service or device so as to increase the amount of data they have for potential trade value. .

What Should Be Done Differently in the Next Research Cycle

At this point, the data marketplace increasingly incentivizes the Quantified Self industry to collect data as a commodity that can be traded commercially in perpetuity, without compensation to those from whom the data was harvested. Left unchecked, it's a form of imperialism, where the data imperialists will benefit, and the wealth derived is not shared with the human beings who provided it. I'm prompted to wonder how this will make the world a better place, as per the mantra of Silicon Valley.

It seems this mantra of the tech world is answered prematurely – the assumption is that it makes the world a better place for the users of their technology. And I agree that it can certainly improve health and wellness. However, in the broader view, hooking users as a requirement to remain competitive in the marketplace arguably has negative impacts for the user. To the extent that it may contribute to inequality – between Data Barons and addicted users – it doesn't seem to make the world the kind of better that I yearn for.



Tech companies measure their success by time-on-site and 7-day-activity metrics, which means that success is defined for them not just by the quality of life they enable for users, but by influencing users to return frequently and stay as long as possible. There is a deep contradiction here. To address this systemic contradiction can be addressed in the next Quantified Self Report Card through a questionnaire that names and explores this contradiction, as well as how it might be addressed. The zeitgeist of making the world a better place is the right spirit and intent. To make it real it is a considerable but worthy challenge.

Reflections: Chelsea Palmer

Concrete Research Process

As the experimental designer- writing the research forms and processing the data structures- this felt like an endless series of re-iterations with tiny adjustments. It was quite satisfying to keep framing the project and its data with minor alterations that made its patterns and conclusions so much clearer, but it was also an exhausting process that often felt like it would never end. However, this whole path was crucial to the exploration process, and I think it was key that I chose to do a full field literature review only *after* our design and collection process was completed. As a lifelong academic nerd, I fear that if I had read all of these articles before the process, I would have inadvertently created a largely derivative and unoriginal methodology. Even though our methodology was quaintly subjective and near self-anthropological, it certainly was unique within the literature that I reviewed afterward.

What I Realized as I Went

We were really addressing a very tiny and specific element of the bigger picture here, but a crucial one-- the interface of the average users' first experience of such a site. This was something I had learned to overlook in my daily life, and the degree of "privacy policy overload" that I began to experience was monumental. However, it did bring a significant element of



mindfulness of my formerly passive downloading of new apps and signing up for new services. The norms that had slowly domesticated me over the past ten years or so- that this is just how it is, I don't have a choice and I might as well not even bother reading the fine print- didn't necessarily disappear but they became much more noticeable.

This deep dive also reinforced for me that it's not just important to recognize how much information is collected about us as users, but where that information ends up, and what it's used for. It wasn't so long ago that our personal health information was kept in paper records, which could not only be purposefully destroyed, but would inevitably be ravaged by time itself. Now, the permanent storage of distributed data is limited only by the capacity of servers-- and these are increasing at a rapid rate to accommodate projected storage needs. The data which is collected today will be readily available for the analytics of tomorrow.

What Should Be Done Differently in the Next Research Cycle

To be honest, I'm mainly hoping to get more hands on deck and distribute the necessary work across a larger group of collaborators. I was disappointed that the logistics and enthusiasm just didn't work out for crowdsourcing the research component, but it became clear that there will need to be at least three dedicated contributors to effectively guide and instigate that process.

Other than the excellent and deeply meaningful reflections of my colleague Rochelle, I wrote this whole excessive document. I was grateful to have feedback and support, but it made me finally realize that something has changed since I was young, and I don't particularly enjoy writing long-form pieces anymore. Additionally, I'm not strong in graphic design or document formatting, so this is an area where we can bring in lots of small contributions from "organic experts" in all the necessary roles for such a massive undertaking.





Appendix C: Personal Post-Script

An Introspective Overview of the Cultural Relevance of Data Ethics by Rochelle Fairfield

Undoubtedly, massive scale data collection and analytics can be an unprecedented boon to humankind. The research questions that can now be posed, explored, and answered opens a world of understanding human and planetary health and well-being that is difficult for even those closest to it to fully grasp or imagine. The quickening of advances in data collection and storage, AI, big data analytics, as well as human-machine interface mean that we can know more about ourselves – individually and collectively – than ever before. And, much of the future value of data collected now is as yet untapped. This is where it gets sticky, since the data collected may be used for purposes that are vaguely if not at all related to the product or service.

It appears that there are implicit assumptions about the relationship between quantity and quality underlying the use of self-metrics. One orienting assumption is that improvements in one's quantitative self-metrics will result in an improvement in one's quality of life. That is, in one's subjective, lived experience: longevity is assumed to make for a better life, for example. These assumptions, whether implicit or explicit, seem to have legitimized the practice of capturing more, rather than fewer, data. Some of the data collected in the quantified self industry is expressly for improving a product or service. Yet, data collection seems to have also morphed into a business practice that isn't for this purpose, but rather to acquire even unrelated data for the purpose of using it to create other potential revenue streams, such as selling it to data brokers, where it's future uses are yet to be tapped.



Human evolution arguably depends to some degree on seizing opportunities and resources. If we can follow the line of thinking that data represents a kind of raw resource, which data analytics then refines into useful information, there are many upsides to how masses of this raw material might be refined to fuel human health and well-being. There is of course a downside as well. Human beings have a tendency to hoard resources, thus creating wealth for some and poverty for others. Our history of colonialism and natural resource extraction across borders is one example, and arguably is still very much practiced. It's plausible, likely even, that this spirit of imperialism will live on in the world of data resource extraction and trade. A new socio-economic category, that of data baron, is currently in the making. We have a chance now to influence and nudge the practices in data collection towards a brighter future, be it Quantified Self or any form of data collection.

Along with these societal level concerns, increasing metrification points to a more existential question: what is a good life? This is a qualitative question that may include quantitative measures as part of the answer, but to mistake measurement for a good life is can easily lead to over commodification of human sentience. Different cultures and world views will have different and layered answers to what makes for a good life. Qualities such as freedom and affiliation mean vastly different things to different peoples. The freedom to liberate some and oppress others, or the insistence that everyone enjoy the same freedoms, for example, suggest that different metrics would be required for the same concept as it is expressed through the multitudes of human expression. Leading us to inquire into the value also of freedom from measurement, and what right humans might have to that freedom.

Self-measurement that takes an objective view of our subjective experience runs the risk of reducing first person human experience to third person facts. The objective lens can, in its



shadow version, lean towards dismissing subjectivity as non-essential or ironically as not valid because it is “too subjective”, which is a violation of the sovereignty of self-hood. Likewise to ‘other’ and dismiss objectivity is a violation of the wholeness of human reality. Truth is, facts matter, and the experience of being matters. Wherever we subjugate one to the other, trouble ensues. Respecting and integrating both subjectivity and objectivity will perhaps provide us the broadest, deepest view of human thriving.

The reach and specificity that current and emerging methods of quantifying the self allows for can help humanity perhaps more than any previous discovery in health, be it germ theory to pharmaceuticals. The power of Quantified Self in a big data era is perhaps comparable in scale to the power nuclear energy over conventional types. The power, reach and impact are several orders of magnitude greater than anything that came before it. Humanity’s rich potential is worth the effort to pre-empt the destructive power of Metrics – especially where self and economic drivers rub against one another – before we have the mess of a melt-down to clean up. To realize the good, we can explore what conditions might be more conducive to realizing more of Metric Culture’s benefits to humanity? What conditions are more likely to result in more harm than good, and how might we steer towards this thriving future? Building ethical best practices into self-quantification is the Human Data Commons’ contribution to this emerging moment in the human story.



Appendix D: Example Research Forms

See attached on next page.



Appendix E: Raw Data Spreadsheets

See attached on next page.



Appendix F: Works Cited

- Acquisti, Agarwal, Bauer, Blum, Breaux, Cranor, Datta, Fienberg, Fong, Jahanian, Jia, Peha, Sandholm, Sadeh, & Sicker (2014). Shaping our national privacy research strategy: A multi-disciplinary perspective. *Request for Information - National Privacy Research Strategy* (October 16, 2014).
- Adams, & Sasse (2001). Privacy in multimedia communications: Protecting users, not just data. In: Blandford A., Vanderdonckt J., Gray P. (eds). *People and Computers XV—Interaction without Frontiers*. Springer, London.
- Alohaly & Takabi (2016). Better privacy indicators: A new approach to quantification of privacy policies. *Symposium on Usable Privacy and Security (SOUPS)*, June 22-24 2016, Denver, Colorado.
- Andrejevic (2013). The big data divide. *International Journal of Communication*, 8.
- Ermakova, Fabian, Baumann, & Krasnova (2014). Privacy policies and users' trust: Does readability matter? (2014). In: Americas Conference on Information Systems. Savannah, USA.
- Haddadi & Brown (2014). Quantified Self and the privacy challenge. *SCL Technology Law Futures Forum*, August 2014.
- Lederer, Dey, & Mankoff (2002). A conceptual model and a metaphor of everyday privacy in ubiquitous computing environments. Report No. UCB/CSD-2-1188. *Group for User Interface Research*, University of California: Berkeley.
- Leibenger, Mollers, Petric, Petric, & Sorge (2016). Privacy challenges in the Quantified Self movement - An EU perspective. *Proceedings on Privacy Enhancing Technologies*, 4.
- Lupton & Michael (2017). 'Depends on who's got the data': Public understandings of personal digital dataveillance. *Surveillance & Society*, 15(2).



- Marotta-Wurgler (2015). Does “notice and choice” disclosure regulation work? An empirical study of privacy policies. NYU Law School dissertation, April 2015.
- McCallister, Grance, & Scarfone (2010). Guide to protecting the confidentiality of Personally Identifiable Information (PII): Recommendations of the National Institute of Standards and Technology. *NIST Special Publication 800-122*.
- Mehlman (2015). If you give a mouse a cookie, it’s going to ask for your Personally Identifiable Information. *Brooklyn Law Review*, 81(1).
- Piwek, Ellis, Andrews, & Joinson (2016). The rise of consumer health wearables: Promises and barriers. *PLoS Med* 13(2).
- Roessler & Mokrosinska (2013). Privacy and social interaction. *Philosophy and Social Criticism*, 39(8).
- Sathyendra, Schaub, Wilson, & Sadeh (2016). Automatic extraction of opt-out choices from privacy policies. *Association for the Advancement of Artificial Intelligence* - <http://www.aaai.org>
- Schaub, Balebako, Durity, & Cranor (2016). A design space for effective privacy notices. *Symposium on Usable Privacy and Security (SOUPS)* July 22-24, 2015, Ottawa, Canada.
- Thierer (2015). The Internet of Things and wearable technology: Addressing privacy and security concerns without derailing innovation. *Richmond Journal of Law & Technology* (21)2.
- Thompson (2011). Shifting boundaries of public and private life. *Theory, Culture & Society*, 28(4).
- Vayena, Mastroianni, & Kahn (2013). Caught in the web: Informed consent for online health research. *Science Translational Medicine Magazine*, 5(23).
- Weiss & Archick (2016). U.S.-EU data privacy: From Safe Harbor to Privacy Shield. *CRS Report: Prepared for Members and Committees of Congress*, May 19, 2016.
- Woo (2006). The right not to be identified: Privacy and anonymity in the interactive media environment. *New Media & Society*, 8(6).

