QUANTIFIED

BIOSENSING TECHNOLOGIES IN EVERYDAY LIFE

EDITED BY DAWN NAFUS

5

Biosensing in Context: Health Privacy in a Connected World

Helen Nissenbaum and Heather Patterson

Introduction

Socio-technical systems that enable communication between "smart" networked devices create a suite of new vulnerabilities for individual users and for society, the precise character of which reflect differences between entrenched and novel social practices. Here we argue that these vulnerabilities implicate privacy by virtue of disrupting settled information flows, and thus warrant a thoughtful consideration of the ends, purposes, and values of the underlying context in which novel information flows occur.[1]

Recent years have seen the emergence of a vast array of personalized self-monitoring tools, the use of which challenges long-entrenched social norms governing personal information sharing. Self-monitoring practices that began with sports and fitness bands and portable medical devices have rapidly expanded into new realms, facilitating caretaking of self and others and forging new inroads into health, wellness, and productivity from cradle to grave. Although these tools present exciting opportunities for raising awareness, building community, and improving efficiency and productivity, they also highlight tensions that animate contemporary privacy debates as individuals and societies discover, adjust to, and resist opportunities afforded by new technologies. A consideration of the precise nature of disruptions introduced by self-tracking technologies will help us determine whether new practices are problematic, and, if so, whether those problems are best approached with an eye toward legislative, policy, or technological solutions.

In this chapter we assess contemporary health self-tracking practices through the lens of Contextual Integrity, an analytical privacy framework that demands a full consideration of the social settings in which novel practices are situated, including the type of information at issue, the identity of the information subjects, senders, and recipients, and the social

norms underlying the context in which new information flows occur. Our approach entails first a description of a novel practice, and second, a normative evaluation of this practice in terms of individual interests, social values, and contextual ends, goals, and purposes. We conclude with an examination of tools at our disposal to bring these systems in line with normative values, including law, policy, and technical design.

A Taxonomy of Health Self-Tracking

To clarify our domain and anticipate terminological ambiguities, we first situate our argument and cases within the larger landscape of health self-tracking. As explained in this book's introduction, the class of health biosensor systems includes great variation in its instances, and along with this, great variation in the terms applied to them. Fitness bands, medical sensors, health apps, smart garments, workplace performance trackers, and home life automation technologies allow for the monitoring of a wide variety of bodily metrics across the lifespan of their wearers and within numerous social contexts. In observing the range of these systems, including their served populations, composition, and aims, we identified several factors that distinguish them:

- *System components* (e.g., hardware, software, and linked services for users and third parties);
- *Device form factors* (e.g., freestanding consoles in the home, wearable clothing and jewelry, patches and tattoos, and implantable or ingestible medical devices);
- *Input modalities* (e.g., automatic sensing, manual self-reporting);
- *Sensor capabilities* (e.g., accelerometers, weight scales, and respirometers);
- *Information types collected* (e.g., distance walked, body mass index, and respiration rate);
- *Aims of the system* (e.g., increasing awareness, deepening self-understanding, building community, automating or facilitating caregiving, enhancing workplace productivity, or managing insurance costs);
- *Actors in the circuits of information flow,* including *Initiators* (e.g., self, care givers), *Data subjects* (e.g., self, children, employees, patients), and *Data recipients* (e.g., self, online or offline friends, family members, colleagues, third-party wellness administrators, human resources personnel, insurance company officers, physicians, public health officials).

Delineating these factors provides axes of systematic differentiation among this burgeoning array of devices and services and helps distinguish between instances on which we focus in this chapter. Ultimately, this will be useful in teasing apart diverse systems and surrounding practices that are relevant to privacy.

Health Self-Tracking and Privacy

At first glance, privacy would seem irrelevant to the domain of self-tracking—tracking *of* one's self *by* one's self. Other normative problems may come to mind: perhaps obsessive perfectionism or intense self-focus—but why privacy? The brief answer is that most self-tracking systems radically interrupt and divert preexisting information flows. According to the definition of *privacy as Contextual Integrity* that we adopt in this chapter, privacy requires the appropriate flow of information, which means flow that meets legitimate expectations. Legitimate expectations, in turn, are characterized by context-specific norms of information flow that not only are entrenched in the practices and conventions of a given context (for example, health care, education, religious practice, etc.), but that also support important ethical and contextual values.

If nothing else, these new technologies disrupt entrenched information flows by virtue of many of the properties we have noted already—namely, close and continuous monitoring of personal information through networked sensory apparatus and self-reporting, and the accumulation and use of this information not only by the data subject, but also by third parties. Additionally, the fact that many trackers are small and unobtrusive encourages ubiquitous wear, and their sensor capabilities and flexible input modalities allow the logging of a full complement of detailed user behaviors, such as when individuals wake, bathe, eat, work, recreate, and sleep. Commonly branded as tools for healthful and positive lifestyles, they tacitly encourage honest and complete engagement. Yet the ecosystem within which these trackers actually function, including interoperable "smart" devices, apps, and services facilitating third-party access to information that may be collected, distributed, assembled, and mined, belies the user's subjective experience of a truly closed system of self-tracking.

These constitute the prima facie infringement of existing informational norms due to radical alterations of the type, frequency, breadth, and depth of information tracked, how it is tracked, and who has access to it. These observations do not answer the question of whether self-tracking systems

violate privacy. They merely indicate why we need to ask the question of whether disruptions are appropriate.

Contextual Integrity

According to the framework of Contextual Integrity, what people care about is not that we should have complete control over information about ourselves, or that no information about us should be shared, but that it should be shared appropriately (Nissenbaum 2009). Contextual Integrity is predicated on the notion that social contexts are an organizing principle of social life, in that people do not act or transact merely as individuals in an undifferentiated social world, but rather as actors, operating in "structured social settings characterized by canonical activities, roles, relationships, norms (or rules), and internal values (goals, ends, purposes)." Context-dependent informational norms embody appropriate information flows by prescribing (and proscribing) what types of information, and about whom, may be transmitted by whom and to whom, and under what constraints. Thus the framework posits the parameters of *actors* (subject, sender, and recipient), *information types*, and *transmission principles*, each ranging over the ontologies that constitute respective social contexts. When actions or practices violate entrenched informational norms, they provoke protest, indignation, or resistance. When actions or practices are in compliance, they respect contextual integrity.

More widely, information technology and digital media have aroused deep concerns over privacy, according to this theory, because they are responsible for massive disruptions in the ways information (or data) *flows*—ways in which it is captured, utilized, and disseminated. Before we can begin to evaluate whether these disruptions should be welcomed or resisted we must clearly understand their nature and sources. Contextual Integrity offers a way to do both—to locate and describe disruptive flows and also to guide assessment in moral terms. A heuristic emerging from the theory suggests key steps: (1) establish a prevailing *context* for the action or practice in question; (2) identify the key *actors* in terms of the context-specific functions or capacities in which they are acting; (3) distinguish which *attributes* (or information types) are in play; and (4) ascertain *principles of transmission* governing information flows.

The heuristic provides a way to compare entrenched practices with those introduced by novel systems: Does the new system introduce different actors into the information flows? Does it offer access to new types of information? Does it shift the terms under which information flows, for example, from "with permission of data subject" to "with payment to service provider"? Detecting differences flags the need for further examination and evaluation. Although the default favors entrenched informational norms, which are presumed to support settled social values and interests, a thorough normative evaluation may favor the patterns that emerge when novel technical systems interject in an information flow.

Conducting a normative evaluation of novel flows involves three layers: one covers *interests*, a second, *general ethical and political values*, and a third, *context-specific ends and values*. The first two are subjects of a large and growing literature on privacy that probes harms and benefits of various systems, devices, and sociotechnical practices they occasion. Beyond interests, it considers whether values are threatened, for example, through unfair discrimination, threats to autonomy, chilling of speech and association, and so forth. Less evident, however, is attention to context-specific aims and values, such as health outcomes and fair allocation of benefits in a medical context, productivity in the workplace, and trust and safety within the home.

To start, context must be counted. Are we considering practices within clinical medical settings, health and fitness communities, places of work, or families? Background context will make a profound difference to how disruptive information flows are experienced by individual users and affect the significance and meaning of these patterns of flow. The perturbations of novel practices on an entrenched system may be positive indeed if they allow for a better realization of the relevant values. Such are the hopes pinned, in the United States and elsewhere, on a contemporary build-out of a health information infrastructure to improve healthcare delivery, medical outcomes, cost efficiency, and public health and well-being. But because health self-tracking is not limited to the healthcare context, a fact emphasized by developers and promoters, it is crucial to locate the various contexts of use in order to track and evaluate relevant information flows and positive and negative impacts on those respective contexts.

Descriptive Evaluation

In what follows, we draw on the Contextual Integrity heuristic to examine and evaluate privacy concerns in the development and use of health self-tracking in one context: that of the workplace. In our view, this case portends a worrisome trend deserving close attention and mitigation. Leading fitness tracking companies may cultivate new markets not only by selling their products and services directly to the public via retailers,

but also by embedding them into existing health and wellness infrastructural ecosystems. For example, an employee might receive a free or discounted fitness tracker from her corporate employer as a benefit of enrolling in the company's workplace wellness program, or on the condition that it be worn for a particular event, such as a fitness competition, or to demonstrate mastery of a particular fitness goal, such as a daily step count average. Its use may even establish eligibility for discounted insurance premiums under the terms of a health-contingent wellness plan. Under a different model, employees might voluntarily use health self-tracking devices to log personal fitness metrics and share this information with colleagues by uploading it to company servers or otherwise making it available for viewing by fellow employees. One firm, for example, aims to track correlations among the fitness, productivity, and happiness of its employees by integrating health, project management, and social interaction data collected in-house and externally (Finley 2013). Are these flows disruptive? And if so, are they morally defensible?

Context. To begin, we place these practices in the overarching context of the workplace. Information flow patterns may vary from one workplace to the next (e.g., Mount Sinai Hospital, the CIA, or Walmart), and may be shaped by physical *layouts* (e.g., corporate offices, open warehouses, or delivery routes) and specific employment *junctures* (e.g., job interviews, performance assessments, or insurance pricing evaluations). However, there are salient common characteristics, including workplace hierarchies and purposes, which yield a significant set of expectations common to most.

Actors. For purposes of our analysis, data *subjects* are tracked company employees. *Recipients* of information generated by tracking systems may include the data subjects themselves, their peers, bosses, and other company officials, and persons associated with the tracking devices and affiliated service providers. Device and service providers also function as data *senders*, as might other entities acting as intermediaries for the purpose fulfilling a program's mandate, such as third-party wellness administrators or insurers who analyze and aggregate user data and send it back to their corporate clients.

Attributes. The type of information in question could vary from case to case, depending on the devices and systems used, as well as the programs and policies of particular workplaces. It might include daily step count or other activity metrics, weight and body mass index (BMI), foods eaten, calories consumed, heart rate, sleep quality, asthma symptoms, mood, and more. This information may be linked with other identifying information, such as demographic data or company department, or with workplace productivity metrics. In general, self-tracking systems provide an unprecedented degree of detail about health information.

Transmission principles. When evaluating new patterns of flow brought about by novel technology-enabled practices, the focus is on alterations to these terms from baseline to novel practices. Because these may be highly variable from case to case, findings will likewise vary. Further, since health self-tracking may introduce new types of information, there may be no baseline practices with which to compare. Taking Fitbit as an example, a subject's activity data are automatically uploaded to the Fitbit server as a condition of using the tool. This principle is different from one governing a subject's self-reporting of mood, energy levels, or alcohol consumption, or of one governing a subject's uploading of data from another self-tracking tool, such as a connected weight scale, glucose monitor, or food tracker. Here, differences hinge on whether data sharing is mandatory or optional, automatic or manual, and also whether it is subject to regular renewal and cross-platform integration.

The character of transmission principles is also determined by employers' general policies, including the terms of employee insurance plans, remuneration programs, and oversight practices concerning health-related lifestyle habits. A leap to automatic transmission of a continuous stream of biometric data for assessment against fitness requirements, for example, would be a huge departure from more settled practices, such as optional, scheduled, in-person biometric screenings with an employer's insurance provider representative. Variation may also be introduced as a function of specific agreements between employers and third-party tracking companies, or between employers and employees. Employers might impose contractual obligations on service providers to make employee information available for analysis. Or, companies might impose "softer" obligations to their employees, communicated as ideals of corporate good citizenship. One company official remarked that health data collected by her company is "not used for anything at this time" and "we don't base any team decisions around it." However, she also noted that if a potential employee is turned off by lifestyle information sharing, "he or she may not be a good cultural fit for the company in the first place" (Nield 2014). This position does not appear to allow for a real choice as to whether to authorize transmission of information. Transmission principles would also govern an employer's policies on sharing employee data with parties outside the company.

Normative Evaluation

Privacy concerns that emerge in the wake of new technologies are often dismissed on grounds that "nothing has really changed" aside from better access, more efficient monitoring, or greater convenience. Such claims may comport with privacy understood in terms of the dichotomy of *private* vs. *public*, or subject *control* of information, or concrete *harms* suffered. The finer lens of Contextual Integrity, however, reveals a different picture. It flags alterations in information flows to which more reductive accounts are blind; it offers a rigorous account of these instead of vague charges of "creepiness," or, worse, of irrationality. In our two cases, but also generally for health self-tracking systems, Contextual Integrity reveals that flows are altered in all three parameter fields—actors, information types, and principles of transmission. Disruptions in any of these parameters, as noted earlier, may be vectors of privacy violation if they fail to meet the criteria of a normative analysis.

Of greatest significance is the new inclusion of employers as recipients of information produced by health self-tracking, which, in our view, inappropriately extends the reach of employers into the lives of employees. (We cannot ignore the irony of retaining the prefix "self" for these flow patterns, as well as those in which systems providers automatically interject their servers as repositories for tracked data.) Unlike prior employment history or performance outcomes, which may even be automatically tracked, health and lifestyle information penetrates into zones formerly reserved for family, friends, or physicians.

One could argue that employers are already expanding their role beyond merely writing a paycheck for work and supervising, managing, and mentoring workers (Henderson 2009). Employers no longer behave as mere *employers*, but additionally assume expansive roles of insurers, benefits providers, behavioral police, and cheerleaders for health, productivity, and happiness. Since employers are paying for health benefits and profit from healthful employees, they have an interest in encouraging healthy lifestyles and attention to fitness and other health metrics. But total involvement by employers may not equally serve employees' interests: continuous tracking may be experienced as illegitimate intrusion into personal life, generating worry and anxiety and preventing relaxation and rest outside working hours. Power differentials between employers and employees may be exacerbated, particularly if policies regarding onward distribution of data are unclear, raising questions about who will get access, under what terms, and with what results.

Researchers and social commentators have challenged incursion such as these into employees' lives. Legal scholar Katherine Stone, for example, calls them collapsed contexts, or "boundaryless workspaces" (Stone 2002), and warns of threats to workplace fairness, equity, and justice. Another legal scholar, Michael Selmi notes, "One of the problems we have in defining the proper space for workplace privacy is that it is no longer clear what work is about, what the boundaries of work are, or even what it means to be an employee. ... It is one thing to give an employer broad dominion over its own workplace but quite another to extend that dominion wherever the employee goes" (Selmi 2006).

Ways in which novel flows affect respective interests constitutes the first layer of analysis prescribed by Contextual Integrity. Presumably, employers' interests are served by producing a harder working, more resilient, and less costly workforce. Yet, for employees, this trend changes the "psychological contract" with the employer, jeopardizing the sanctity of spaces for relaxation, reflection, and experimentation (Stone 2002). It also expands zones in which employers may exercise power over their employees.

Especially worrisome is the prospect of workplace discrimination if an employer infers that an employee has (or may develop) an impairment that limits her ability to work, or increases health-care costs. In a recent study of contextual expectations of health information flows, Fitbit users strongly resisted employer access to health and wellness data (Patterson forthcoming). One research participant objected, for example: "I don't want an employer or a potential employer to go and find all my diabetes, all my blood pressure, blood glucose, and weight, and all this other medical information, and then say 'This guy's going to drive our healthcare [costs] up.' And so, I don't even get the interview because [a potential employer] has just tremendous insight into my health before he even contacts me." Further, research participants were keenly aware of associated harms, such as disadvantages in hiring or promotion processes. For example, they flagged that information about weight or eating habits may signal an inability to exert the kind of discipline and self-control that is valued in the workplace; that information about sleep cycles may signal poor productivity; that information about moods may signal depression or general instability or unreliability; and that information about family histories of diseases like cancer, diabetes, or heart disease may signal increased insurance costs and place them at risk of dismissal.

These fears are not irrational. In the early 1990s, Indiana's Best Lock Corporation, whose policy prohibited the use of alcohol, drugs, and

tobacco both at and away from work, attempted to fire an employee for disclosing that he had once consumed alcohol at a bar with friends (*Best Lock Corp. v. Review Bd. 1991*). In 2007, Scotts Miracle Grow successfully dismissed a newly hired lawn-care technician for violating the company's nicotine-abstinence policy (*Rodrigues v. The Scotts Company, LLC et al. 2009*). And in the past few years, Methodist Hospital System, Baylor Health Care System, and Citizens Medical Center have announced policies against hiring applicants who use nicotine or whose body mass indices (BMIs) indicate that they are obese (Roberts 2014).

When considering impacts on ethical and political values, as directed by Contextual Integrity's second layer of analysis, questions of fairness are raised by employers evaluating workers on the basis of health, fitness, and lifestyle choices instead of exclusively on the quality of their work. A traditional understanding of the scope of workplace performance evaluation may encompass, for example, an assessment of an employee's skills, work output, productivity, efficiency, intelligence, originality, creativity, speed, reliability, consistency, honesty, meticulousness, and ability to work well with others. The introduction of health criteria into this space extends the image of the idealized "good" employee to someone who is fit, trim, and even-tempered, who sleeps well, and has few or no vices outside of work. Much as Henry Ford instructed dozens of investigators in his Social Welfare Department to spontaneously visit the homes of employees and flag behaviors that violated company policies, such as gambling, drinking, taking on lodgers, and engaging in sexual relations outside of marriage (Maltby 2008), today's employers implement strict policies related to their actual and future employees' offsite behaviors (Roberts 2014).

Heightened scrutiny and uncertainty may lead to what James Hoopes calls "management by stress" (Hoopes 2005) and what University of Kansas Professor Jerome E. Dobson refers to, with respect to GPS tracking of employees, as "geoslavery" (Dobson and Fisher 2003). Professor Jeffrey Rosen cites sociologist Erving Goffman for the proposition that job tensions are increased when employees must perform under constant scrutiny (Rosen 2001), and notes that workers "experience a dignitary injury when they are treated like the inhabitants of the Panopticon" (Rosen 2001). These outcomes do not merely spell harm, but threaten values to which the majority of Americans subscribe, such as individual autonomy, respect for persons, just deserts, and fair treatment.

Even if workplace wellness programs are framed as voluntary, employees may experience economic pressure to participate, social pressure to conform, and diminishing privacy expectations. Law professor Scott Peppet notes that when disclosure of personal information is economically attractive to employers, as well as inexpensive, easy, and common, pressure builds to disclose. Those who resist may be seen as withholding negative information from the community, and thus be stigmatized and penalized (Peppet 2011). Consequences of nonparticipation in health-tracking programs can be serious: by charging unequal health insurance rates to differently situated individuals unless they achieve a particular set of health outcomes, wellness programs place an extra burden on the poor, the sick, and those genetically predisposed to obesity, heart disease, diabetes, and alcoholism. Consequently, those in greatest need are likely to be worst served by such developments.

The third layer of Contextual Integrity analysis looks for impacts on ends and values of the workplace context. One consideration is organizational stability, achieved through the cultivation of a well-trained, secure, and established workforce; trust in the fair allocation of rewards, reflecting workplace effort, education, training, ability, and interest; and structural arrangements that support workplace cooperation. Social psychologist Roderick Kramer, for example, observes that small gestures, such as the elimination of storeroom locks or time clocks, or the presence of policies encouraging employees to borrow company equipment on the honor system, may signal to employees that they are trusted by their managers, and in turn facilitate an expectation of reciprocal cooperation, "creating a shared common knowledge of the ability of the players to reach cooperative outcomes" (Kramer 2006).

Individuals are adept at "neutralizing" undesirable surveillance by avoiding situations in which they will be monitored, by masking or distorting their identities, and by plainly refusing to participate in data collection activities (Kramer 2006). Gary T. Marx notes that workers whose productivity is evaluated, for example, by the number of typed keystrokes they produce in a fixed period of time may resort to distorting their true efficiency by pressing a single key for several minutes at the end of the tracked window, perhaps obfuscating their output in order to preserve autonomy or to bring their scores into alignment with performance expectations (Marx 2003). The overseer of a state health policy center recounts that students forced to participate in annual fitness "weigh-ins" thwarted intrusive surveillance by wearing ankle weights under their jeans (Hoffman 2015). Defensive strategies such as these may undermine the social and economic success that the programs are designed to promote. Employees subjected to real-time health monitoring may similarly

find ways to subvert surveillance in a cat-and-mouse struggle to achieve a discounted annual health insurance premium, and, in so doing, adopt an oppositional stance within the workplace.

These forecasts of dysfunctional work environments resulting from the insidious surveillance of health tracking data, while merely speculative, are drawn from insights into the outcomes of other forms of close worker surveillance. One final observation concerns not so much workplace as workforce. An efficient distribution of human resources would match individuals with work to which they are best suited. The United States prohibits various forms of workplace discrimination, such as discrimination based on gender and race, for primarily ethical reasons. But there is also an efficiency argument to consider. Prejudicial hiring means that the most qualified candidates will not necessarily be chosen. Similarly, short-term prejudice against well-qualified candidates based on health factors might threaten the overall quality of the workforce. Although we are not in a position to validate this point with economic data, Contextual Integrity would suggest examining this thesis as an aspect of the third layer of analysis concerned with contextual goals and values.

Mechanisms for Protecting Privacy

We have focused primarily on norms of flow supported by a context-based analyses of disruptions posed by health self-tracking systems. In this section the question is practical: what means do we have for enforcing, shaping, or merely encouraging the adoption of practices that comply with these norms? We fear that individuals have unrealistic expectations of the legal system's ability to protect privacy in health-tracking data. They may incorrectly believe self-generated information to be medical in nature, and therefore covered by health privacy laws. Let us take a hard look at existing offerings, and consider gaps and vulnerabilities as well as hopeful future directions for closing these gaps.

Architecture

Any design plan must consider what information is collected, how it is collected, how it is used, to whom it flows under which circumstances, how it is to be stored and maintained, and under what circumstances it is to be destroyed. Design choices, therefore, reveal a privacy and security footprint for any system (Lessig 2006). By the same logic, a consideration of these choices allows the creators of a system to design for privacy and security. Health self-tracking systems may incorporate several

junctures at which data are collected, processed, and disseminated. For example, under one model, after information is collected automatically from fitness bands, it is transferred to a company's cloud server where it is integrated with information from other devices and services. Data are algorithmically analyzed according to various health metrics and then is delivered to personalized dashboards on the service provider's website, or on the user's smartphone app. These data can then be pushed to the end-user's friends via social media feeds, pulled by software developers and their customers via application programming interfaces (APIs), and shared with affiliated third-parties or other business partners, such as employers or insurance companies. Designs that seem inevitable to non-expert users are rife with choices that could have been made differently: for example, whether gathered information first pauses on a user's system so that she may decide whether to allow or restrict onward flow, and whether she may then engage with permissions that are built into a system in order to customize access to others at her discretion. We mention information collection and transfer points to highlight that each juncture presents an opportunity to apply privacy-enhancing design choices. If we want architecture to embody these constraints, we must first know what the constraints are, and we should have the means of holding designers accountable for their design choices.

Law

Much of United States privacy law has developed in relation to concerns about specific technological advances that surprised and worried the public. When Samuel Warren and Louis Brandeis wrote their now-famous 1890 *Harvard Law Review* article advocating for the legal codification of a "right to be let alone," they were reacting to a particular confluence of new inventions and business models that concomitantly enabled new information flows: flash photography, printed newspapers, and other machines allowed for the easy collection and rapid circulation of visual images—and accompanying crops of unseemly gossip—that threatened to make real the ominous fear that "what is whispered in the closet shall be proclaimed from the house-tops" (Warren and Brandeis 1890).

Beginning in 1903, a number of privacy laws responsive to the indignities that Warren and Brandeis foreshadow emerged from the courts and legislatures. However, unlike its European counterparts, the United States Congress has repeatedly declined to pass omnibus federal privacy legislation. Rather, it has enacted a suite of sectoral laws applying to information circulating within particular contexts such as education,

health care, and finance. Additionally, administrative agencies such as the Federal Trade Commission (FTC) and the National Telecommunications Information Association (NTIA) codify and enforce regulations on relatively cabined aspects of social life, such as consumer protections or telecommunications.

Although often criticized as insufficiently protective of privacy, the United States' sectoral regime implicitly acknowledges that privacy hinges on contextual factors, imagined and valued by persons in somewhat different ways, at different times, against more powerful norms-governed backdrops. Indeed, individuals often organize their social behaviors to avoid violating norms underlying these contexts. Thus, even for identical data—say, one's weight or blood pressure or heart rate—selective sharing and withholding of information is not only acceptable, it is often socially obligatory. Health confidences made to one's family, friends, or physician, for example, would be odd and perhaps unwelcome if made to one's employer, first date, or PTA board, and could result in discrimination, rejection, or ostracism.

As a general matter, regulation of health self-tracking is problematic because much of the information collected and processed by commercial sensors and app companies is closely aligned with, and sometimes identical to, data collected in the traditional medical context, but its privacy and security are not specifically subject to privacy regulations relating to health care. One might imagine, for example, that privacy rules associated with the Health Insurance Portability and Accountability Act (HIPAA) would be a promising avenue for protecting privacy in relation to health self-tracking data. However, health self-tracking information does not usually fall under the purview of HIPAA because the law is limited to discrete health-care relationships, rather than health information. Where physicians or insurance plans are subject to restrictions regarding storage and distribution of their patients' or customers' health self-tracking data, commercial actors and others who hold the same data are not. In the cases of health information in the employment context or in home life, protections afforded by HIPAA do not necessarily apply.

Employees do enjoy some limited protection through a patchwork of legal rules designed to mitigate employment discrimination and coercion. For example, in the case of workplace wellness programs that qualify as HIPAA-covered health plans, HIPAA's privacy provisions allow employers access to only "summary health information" without employees' written consent. Under HIPAA's anti-discrimination provisions, employers are required to limit the size of the financial incentive they offer to employees

for participating in wellness programs, and to offer a reasonable alternative to employees who do not qualify for a program. The consent provision in these instances is laudable, but in practice it weakens HIPAA restrictions. Under one insurance company's employee wellness program, for example, an Authorization and Purchase Form for the Fitbit pedometer asks employees to sign a waiver consenting to their step counts being viewed by their human resources department, and to avow honesty in earning step counts and performing walking activities (United Healthcare 2012). Further, outside the rubric of bona fide wellness plans, the voluntary disclosures to companies, such as those undertaken in the spirit of workplace camaraderie discussed earlier, are not covered by HIPAA.

The Americans with Disabilities Act (ADA) and the Genetic Information Nondiscrimination Act (GINA) offer protections against workplace health-based discrimination, but these are also only partial. The former limits circumstances under which employees may be required to submit to medical examinations or meet a health standard in a health-contingent plan; the latter prohibits discrimination based on genetic information by group insurance plans and by employers. However, ADA provisions are only triggered if a disability is in play. Commentators generally believe that the ADA offers little recourse against discrimination based on lifestyle factors (Roberts 2014). Employee information protected by GINA is limited to genetic information and does not include information about sex, age, race, or ethnicity that is not acquired via a genetic test.

Focusing on federal law covering the employment context, we find that employees currently have few rights against private sector employers who engage in various forms of monitoring or tracking. Private sector employers are permitted to search and engage in widespread surveillance of employees, including testing for drug use, tracking employees' phone calls and texts, monitoring employees' whereabouts via GPS or RFID tags implanted in clothing or badges, and subjecting applicants or employees to honesty and personality tests. A number of states have "lifestyle statutes" that protect the rights of employees to engage in lawful activities outside of the workplace during nonworking hours, such as smoking and drinking alcohol. These statutes, however, are not absolute and allow employers to take action if the behavior in question conflicts directly with essential business-related interests, or violates a legal or contractual agreement with the employee. The common law privacy torts may offer some protections to employees but are similarly limited in scope. The tort of Intrusion Upon Seclusion, for example, provides a cause of action where one's solitude or seclusion is intentionally intruded upon in

one's private affairs or concerns in a manner that is "highly offensive to a reasonable person" (American Law Institute 1977). With rare exception, narrow interpretations of "private affairs or concern" provides an employee with little recourse to monitoring that can plausibly be related to employment, even if undertaken outside of the workplace, and any careful "reasonable person" analysis must take into account changing social practices ushered in by new technologies.

Policy

As we have seen, existing legal protection covers only thin slices of the domain of health-tracking systems. Outside the scope of these legal regimes, the practices we discuss fall within the regulatory framework of commercial actors whose business involves collecting and using personal information. In the United States, this extensive landscape ranges over companies that collect information in the process of providing other goods and services to those whose business *is* information. Under the dominant regime, self-regulation, companies declare self-fashioned information practices through privacy policies and, in principle, are held accountable to these by government agencies, primarily the FTC.

Regulatory bodies and advocacy organizations have urged companies to base their voluntary policies on the Fair Information Practice Principles (FIPPs), a set of principles first articulated in the 1970s by a Committee of the U.S. Department of Health, Education, and Welfare and, virtually simultaneously, adopted by the Organization for Economic Co-operation and Development (OECD) (Regan 1995). Developed in response to deep worries over the growing use of computerized database technologies, the FIPPs have been a cornerstone of government and commercial privacy regimes, and have been formulated to provide clear notice to individuals regarding the collection, use, dissemination, and maintenance of data, to articulate the purpose or purposes for which this information is collected and used, to retain only data relevant and necessary to accomplish those purposes, and to seek consent for these practices or departures from them. Further, data quality and security provisions should ensure that stored information is accurate, relevant, timely, and complete, as well as protected from loss, unauthorized access or use, or disclosures through appropriate security.

Companies wishing to present themselves as conscientious actors generally adopt the rhetorical stance of presenting privacy policies clearly and implying that individuals are free to consent to them, or not. The reality, as critics have pointed out, is a far cry from the ideal. Even when policies are written with care and with fidelity to FIP principles, they are notoriously long and difficult to understand. Ample research shows that most people do not read privacy policies and that when they do, they do not grasp them. Often, policies are deliberately written broadly enough to encompass a vast range of behaviors—present and future—in order to protect companies against legal liability through inadvertent violations.

A major source of weakness is that the FIPPs offer primarily procedural and not substantive protections. The FIPPs are envisioned to level the playing field between more and less powerful actors by codifying procedural rules of fairness and by providing a measure of control to individuals through consent mechanisms. In practice, however, companies may be virtually unrestricted in the practices they follow, so long as they declare them in those same policies that no one reads or understands. This opens a great loophole, important to mention here, but discussed in greater detail elsewhere (Nissenbaum 2015). An alternative regime for regulation suggested by Contextual Integrity is through substantive rules derived from ideal contextual informational norms.

Regulators are aware of self-regulatory shortcomings and have recently contemplated how best to approach the protection of consumer data, including self-generated health information. In May 2014, for example, the FTC held a seminar examining the collection, use, and distribution of health data generated by consumers and corporations outside of the traditional health-care context (Federal Trade Commission 2014). And the White House's 2015 Consumer Privacy Bill of Rights discussion draft is prescient in recognizing that consumers often engage with technology differently as a function of the social contours of a particular business sector or environment (White House 2015). However, respecting context requires an awareness of what consumers expect from a particular social context, and of what is at stake when information flows are disrupted (Nissenbaum 2015).

Future Work

As health data move into employment and other spaces, previously unheeded relationships will need to be regulated by a suite of complementary legal, policy, and architectural formations. Whether externally imposed or voluntarily adopted, new privacy-protective measures will benefit from a clear articulation of the nexus between data flows and contextual values. Contextual Integrity provides a procedural roadmap, but data are needed to inform fuller analyses that can translate to concrete and appropriately nuanced architectural, policy, and legal solutions.

Fundamentally, a commitment to contextual privacy demands a commitment to radical transparency. The onus here is on companies to self-disclose, and on legislatures to require self-disclosure. Transparency is vital because it is impossible to critically evaluate a new program without fully understanding what consumer information is collected and how it is used. Thus, "radical transparency" in our view entails explicit point-to-point flows of information that reveal precisely what type of information is disseminated, to which parties, and under what conditions. Technologists, social scientists, philosophers, policy analysts, and legislators each have a role to play in defining and shaping the contours of these developing systems. Particularly needed are careful audits of information collections and flows, coupled with social science research that delves deeply into users' expectations for (and practices with) new technologies. When combined with a clear identification and articulation of contextually grounded values, this resulting knowledge base should enable the identification of gaps between actual and desired legal protections, and from there, the enactment of laws regarding which informational transactions should proceed, and under what conditions.

To illustrate: in the situation of health self-tracking data making its way into the employment context, the patchwork of legal regulations on employers are so sparse—and employees so thinly protected—that new federal privacy regulations of actors distributing and receiving data may be warranted. This could take many forms, including developing a broad federal statute protecting employee privacy rights with respect to surveillance, of which health surveillance is but one type; extending HIPAA to cover employers, per se, regardless of the administrative path (e.g., wellness plans tied to insurance plans) through which health information is collected; or even by removing health information from the workplace entirely. Defining the character of these types of safeguards requires a much closer look at all facets of the information flows we have identified in this chapter, bolstered with a careful analysis of how these restrictions operate in practice.

One reasonable approach may be to segregate identifiable data that users provide for a given purpose, such that it is available for use in that context but unavailable for migration to other spheres. For example, early empirical work suggests that individual and familial genetic profiles, mental health characteristics, sexual and reproductive health histories, and lists of current medication would be strong candidates for mandatory segregation in the employment realm, hedging against invasive health surveillance practices that threaten the mutual respect

between worker and employer that underlies a cooperative workplace. In practice, this would place most health information flows to employers off limits by default, and perhaps render them inalienable even with employee consent.

But it may be the case that users sanction certain types of information flows if conditions are met, such as nearly absolute assurances of data de-identification and encryption, or a strongly restricted recipient list, or a demonstrably and solely prosocial use of the data. For example, some individuals may choose to permit the distribution of data for research purposes when it is closely tied to the topic of the research study (e.g., genomic data for gene sequencing research). It should not be assumed that in doing so they are giving either explicit or tacit permission to use the data for generalized research purposes. Similarly, individuals may tolerate or even embrace certain data flows to medical personnel such as physicians, but only under the conditions that the information be held in the strictest of confidence and be distributed only when it may reasonably help another medical professional assess a condition that the data subject wants evaluated.

What conditions underlie trust in health information flows in an employment context? To fully grasp these issues, certain information must be acquired, such as how employees who are subject to such scrutiny respond to it in practice, for better and for worse. What are employee fears? What would an ideal information-flow system look like to workers in different types of employment situations? Under what conditions, if any, would employees and colleagues be privy to which types of health data, and why? But before these questions can be answered, we are sorely lacking fundamental factual details:

- What, in practice, are the myriad channels through which health data are currently introduced into the workplace, and how, precisely, do different types of tracking technologies alter the nature or number of these channels?
- What is the exact nature of the data that get introduced through each channel? What, if any, of this information is combined with additional data to reveal new insights about employees? If this occurs, what additional data are accessed and analyzed, and what insights are revealed?
- How and by whom are various categories of information stored, analyzed, and shared within a variety of workplaces? With what personnel? And which informal, and thus potentially unexpected, information dissemination routes are active in various types of workplaces—for

instance, observation, small talk, and gossip—and what are the ramifications of these for employees?

Without these facts, information-sharing practices are subject to ongoing speculation, as is a full accounting of their positive and disruptive effects.

Note that in laying out the need for future empirical work, we are not suggesting that either wholesale resistance to or adoption of new self-tracking technologies is the correct course of action—but rather, that solutions will benefit more from specific knowledge than from broad sketches. We encourage entities that conceive, design, and deploy domestic health self-tracking systems to consider relevant cultural and social contexts in order to more effectively incorporate notions of contextually appropriate information flow into privacy-protective legal and policy frameworks.

Summary

In this chapter we have argued that novel information flows brought about by new health self-tracking practices are best evaluated according to the ends, purposes, and values of the contexts in which they are embedded. Health self-tracking practices may, for example, heighten power disparities between data subjects and recipients that undermine the internal stability of spheres such as employment, where autonomy and freedom from surveillance are necessary to create a productive and harmonious workforce. In others not discussed here, such as the domestic sphere, they may introduce new and unexpected secondary surveillance by third parties, potentially unsettling caregivers in a sensitive environment where trust and security are paramount. Self-tracking tools and practices thus provide intriguing opportunities to explore the philosophical roots of information flow conflicts that occur in relation to specific actors at particular times in a variety of places: within our bodies, our workplaces, our homes, and more.

As society collectively adjusts to value evolutions brought about by new technologies, the proper roles of technologies, laws, and policies may also shift. Architectural, legal, and policy solutions governing the entities collecting and distributing information should be considered in light of, and developed to complement, regulations of entities receiving it. One site of contestation may warrant broad changes in federal regulatory oversight; another may be more effectively addressed through the adoption of new technological or policy practices that maximize information flow transparency and thus better enable individuals to operationalize and enrich values in a particular sphere. Empirical work is essential to the creation of solutions that serve the needs of individuals and society and better enable underlying social norms to weather rapid technological advancements into the future.

Note

1. The authors contributed equally to this work and are listed in alphabetical order.

References

American Law Institute. 1977. 652B Intrusion Upon Seclusion. Restatement of the Law, Second, Torts.

Best Lock Corp. v. Review Bd., 572 N.E.2d 520 (Ind. Ct. App. 1991).

Dobson, Jerome E., and Peter F. Fisher. 2003. "Geoslavery." *IEEE Technology and Society Magazine* (Spring): 47–52.

Federal Trade Commission. 2014. "Spring Privacy Series: Consumer Generated and Controlled Health Data." https://www.ftc.gov/news-events/events-calendar/2014/05/spring-privacy-series-consumer-generated-controlled-health-data. Accessed August 10, 2015.

Finley, Klint. 2013. "What if Your Boss Tracked Your Sleep, Diet, and Exercise?" *Wired UK*, April 18. http://www.wired.co.uk/news/archive/2013-04/18/quantified-work-citizen. Accessed August 10, 2015.

Henderson, M. Todd. 2009. "The Nanny Corporation." *University of Chicago Law Review* 76: 1517–1611.

Hoffman, Jan. 2015. "'Body' Report Cards Aren't Influencing Arkansas Teenagers." *New York Times*, August 10. http://well.blogs.nytimes.com/2015/08/10/body-report-cards-arent-influencing-arkansas-teenagers/?_r=0. Accessed August 10, 2015.

Hoopes, James. 2005. "The Dehumanized Employee" *CIO Magazine* http://www.cio.com.au/article/165305/dehumanized_employee/. Accessed August 8, 2015.

Kramer, Roderick. 2006. "Social Capital in the Workplace." In *Social Psychology of the Workplace*, ed. Shane R. Thye and Edward J. Lawler, 1–30. Oxford, UK: Elsevier JAI Press.

Lessig, Lawrence. 2006. *Code v2*. New York: Basic Books.

Maltby, Lewis. 2008. "Whose Life Is It Anyway?: Employer Control of Off-Duty Smoking and Individual Autonomy." *William Mitchell Law Review* 34 (4): 1639–1649.

Marx, Gary T. 2003. "A Tack in the Shoe: Neutralizing and Resisting the New Surveillance." *Journal of Social Issues* 59 (2): 369–390.

Nield, David. 2014. "In Corporate Wellness Programs, Wearables Take a Step Forward." *Fortune*, September 3. http://fortune.com/2014/04/15/in-corporate-wellness-programs-wearables-take-a-step-forward//. Accessed August 10, 2015.

Nissenbaum, Helen. 2009. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Palo Alto, CA: Stanford University Press.

Nissenbaum, Helen. 2015. "Respect for Context: Fulfilling the Promise of the White House Report." In *Visions of Privacy in the Modern Age*, ed. Marc Rotenberg, Jeramie Scott, and Julia Horwitz, 152–164. New York, NY: The New Press.

Patterson, Heather. Forthcoming. "Contextual Expectations of Privacy in Self-Generated Health Data." Working Paper. New York University.

Peppet, Scott. 2011. "Unraveling Privacy: The Personal Prospectus and the Threat of a Full Disclosure." *Northwestern University Law Review* 105 (3–4): 1153–1203.

Regan, Priscilla M. 1995. *Legislating Privacy: Technology, Social Values, and Public Policy*. Chapel Hill, NC: The University of North Carolina Press.

Roberts, Jessica L. 2014. "Healthism and the Law of Employment Discrimination." *Iowa Law Review* 99:573–635.

Rodrigues v. The Scotts Company, LLC et al. (639 F. Supp. 2d 131 2009).

Rosen, Jeffrey. 2001. *The Unwanted Gaze: The Destruction of Privacy in America*. New York, NY: Random House.

Selmi, Michael. 2006. "Privacy for the Working Class: Public Work and Private Lives." *Louisiana Law Review* 66 (4): 1046–1056.

Stone, Katherine V. W. 2002. "Employee Representation in the Boundaryless Workplace." *Chicago-Kent Law Review* 77 (2): 773–819.

United Healthcare. 2012. "Healthy Directions 2012 Employee Fitness and Wellness Program, fitbit [sic]." http://www.airbornemx.com/benefits/HEALTH%20AND%20WELLNESS%20AMES%20PROGRAM%20detailed%20communications%20to%20employee%20-%20HRA,%20fitbit,%20and%20health%20rewards.doc. Accessed October 25, 2014.

Warren, Louis D., and Samuel D. Brandeis. 1890. "The Right to Privacy." *Harvard Law Review* 4 (5): 193–220.

White House. 2015. "Administration Discussion Draft: Consumer Privacy Bill of Rights Act of 2015." https://www.whitehouse.gov/sites/default/files/omb/legislative/letters/cpbr-act-of-2015-discussion-draft.pdf. Accessed February 28, 2015.

6

Disruption and the Political Economy of Biosensor Data

Brittany Fiore-Gartland and Gina Neff

Science and Technology Studies has long held that the frames and definitions designers give to new tools matter enormously for how users initially receive and ultimately modify those tools. Discourses are powerful forces in technology design, shaping, for instance, how gender and racial inequalities get designed into technologies (Suchman 2002). The startups working in biosensing and self-tracking present a case to examine the role that power plays in the discursive process of framing new technologies. One frame often used for defining new data tools and services includes their abilities for "disruption," or the perceived ability of technologies to upend the status quo of power within established industries and social institutions. In this chapter we present findings from our research in the startup environment in the relatively less-regulated consumer wellness field and the more closely regulated field of mobile medical applications. We use two cases of health data innovation to present possibilities for scholars and practitioners to think about both the processes and discourses of disruption, and how these discourses might affect the design and use of new technologies. Our goal here is not to make normative or evaluative judgments about the roles that disruption discourses play in society. We hope to show that disruption discourses limit how people imagine technologies could bridge existing social contexts and categories. Disruption limits such vision by overlooking the distinct roles for and relationships around data across contexts. People can have different expectations for data within and across different social institutions (Fiore-Gartland and Neff 2015). Social institutions, too, produce the tools and methods for making data intelligible across different contexts. However, the framework of disruption at best ignores social institutions and at worst maligns them. These ways of talking about disruption help to reproduce existing institutional power, even as people use the term