CLEARLY OPAQUE

PRIVACY RISKS OF THE INTERNET OF THINGS

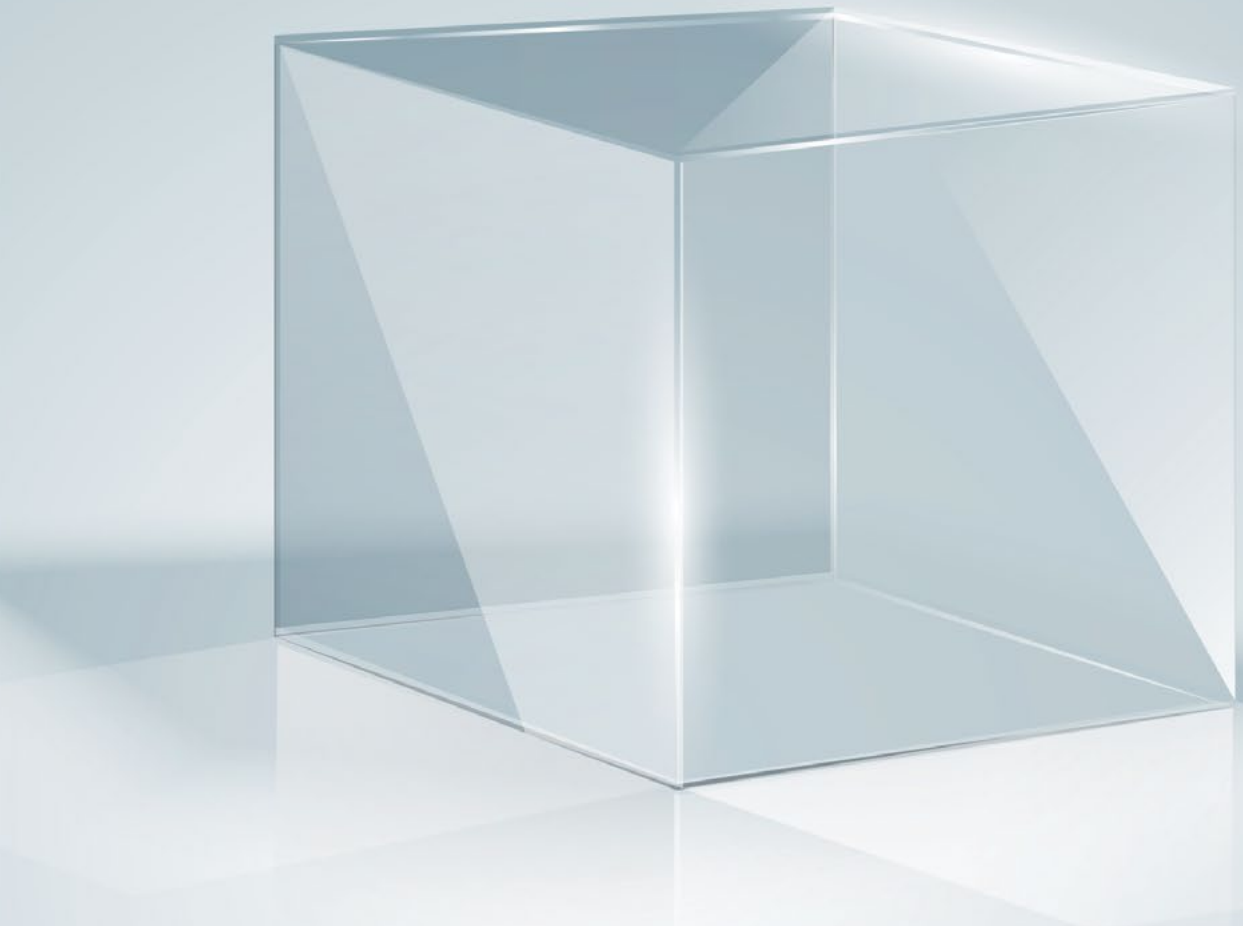

THE INTERNET OF THINGS
PRIVACY FORUM

May 2018

EXECUTIVE SUMMARY

Key IoT privacy risks and issues

The IoT will expand the data collection practices of the online world to the offline world.

— The IoT will enable and normalize preference and behavior tracking in the offline world. This is a significant qualitative shift, and a key reason to evaluate these technologies for their social impact and effect on historical methods of privacy preservation. The very notion of an offline world may begin to decline.

The IoT portends a diminishment of private spaces.

— The scale and proximity of sensors being introduced will make it harder to find reserve and solitude. The IoT will make it easier to identify people in public and private spaces.

The IoT will encroach upon emotional and bodily privacy.

— The proximity of IoT technologies will allow third parties to collect our emotional states over long periods of time. Our emotional and inner life will become more transparent to data collecting organizations.

Given the likelihood of ubiquitous data collection throughout the human environment, the notion of privacy invasion may decompose; more so as people's expectation of being monitored increases.

— Much of consumer IoT is predicated on *inviting* these devices into our lives. The ability to know who is observing us in our private spaces may cease to exist. The IoT will hasten the disintegration of the 'reasonable expectation of privacy' standard as people become more generally aware of smart devices in their environments.

When IoT devices fade into the background or look like familiar things, we can be duped by them, and lulled into revealing more information than we might otherwise. Connected devices are designed to be unobtrusive, so people can forget that there are monitoring devices in their environment.

IoT devices challenge, cross and destabilize boundaries, as well as people's ability to manage them.

— The home is in danger of becoming a 'glass house,' transparent to the makers of smart home products. And, IoT devices blur regulatory boundaries – sectoral privacy governance becomes muddled as a result.

As more and more products are released with IoT-like features, there will be an "erosion of choice" for consumers – less of an ability to not have Things in their environment monitor them.

Market shifts towards 'smart' features that are intentionally unobtrusive lead to less understanding of data collection, and less ability to decline those features.

The IoT retrenches the surveillance society, further commodifies people, and exposes them to manipulation.

The IoT makes gaining meaningful consent more difficult.

The IoT is in tension with the principle of Transparency.

The IoT threatens the Participation rights embedded in the US Fair Information Practice Principles and the EU General Data Protection Regulation.

IoT devices are not neutral; they are constructed with a commercial logic encouraging us to share. The IoT embraces and extends the logic of social media – intentional disclosure, social participation, and continued investment in interaction.

The IoT will have an impact on children, and therefore give parents additional privacy management duties.

— Children today will become adults in a world where ubiquitous monitoring by an unknown number of parties will be business as usual.

Emerging Frameworks and Strategies to address IoT Privacy

Having broad non-specialist social conversations about data (use, collection, effects, socioeconomic dimensions) is essential to help the populace understand the technological changes around them. Privacy norms must evolve alongside connected devices – discussion is essential for this.

Human-Computer Interaction (HCI) and Identity Management (IDM) are two of the most promising fields for privacy strategies for IoT.

A useful design strategy is the 'least surprise principle' – don't surprise users with data collection and use practices. Analyze the informational norms of personal data collection, use and sharing in given contexts.

Give people the ability to do fine-grained *selective sharing* of the data collected by IoT devices.

Three major headings for emerging frameworks and strategies to address IoT privacy:

— User Control and Management
— Notification
— Governance

User Control and Management Strategies

— Pre-Collection
- Data Minimization – only collect data for current, needed uses; do not collect for future as-yet-unknown uses
- Build in Do Not Collect 'Switches' (e.g., mute buttons or software toggles)
- Build in wake words and manual activation for data collection, versus the truly always-on
- Perform Privacy Impact Assessments to holistically understand what your company is collecting and what would happen if there was a breach

— Post-Collection
- Make it easy for people to delete their data
- Make it easy to withdraw consent
- Encrypt everything to the maximum degree possible
- IoT data should not be published on social media or indexed by search engines by default – users must review and decide before publishing
- Raw data should exist for the shortest time possible

— Identity Management
- Design strategies:
  > Unlinkability – build systems that can sever the links between users'

activities on different devices or apps
> Unobservability – build or use intermediary systems that are blind to user activity
- Give people the option for pseudonymous or anonymous guest use
- Design systems that reflect the sensitivity of being able to identify people
- Use *selective sharing* as a design principle
> Design for fine-grained control of data use and sharing
> Make it easy to "Share with *this* person but not *that* person"
- Create dashboards for users to see, understand and control the data that's been collected about them
- Design easy ways to separate different people's use of devices from one another

— Notification Strategies
- Timing has an impact on privacy notice effectiveness.
- Emerging privacy notice types:
> Just-in-time
> Periodic
> Context-dependent
> Layered
- Test people's comprehension of privacy policies
- Researchers are exploring privacy notification automation:
> Automated learning and setting of privacy preferences
> Nudges to encourage users to think about their privacy settings
> IoT devices advertising their presence when users enter a space

— Governance Strategies
- Creation of baseline, omnibus privacy laws for US
- Regulations restricting IoT data from certain uses
- Regulator guidance on acceptability of privacy policy language and innovation
- Requirement to test privacy policies for user comprehension
- Expansion of "personally-identifiable information" to include sensor data in the US
- Policymaker discussions of the collapse of the 'reasonable expectation of privacy' standard
- Greater use of the 'precautionary principle' in IoT privacy regulation
- More technologists embedded with policymakers
- Trusted IoT labels and certification schemes

ABOUT THE AUTHORS/ ACKNOWLEDGEMENTS

DR. GILAD ROSNER

Dr. Gilad Rosner is a privacy and information policy researcher and the founder of the non-profit Internet of Things Privacy Forum, whose mission is to produce guidance, analysis and best practices to help industry and government reduce privacy risk and innovate responsibly in the domain of connected devices. Gilad's broader work focuses on identity management, US & EU privacy and data protection regimes, consumer protection, and public policy. His research has been used by the UK House of Commons Science and Technology Committee report on the Responsible Use of Data and he has contributed directly to US state legislation on law enforcement access to location data, access to digital assets upon death, and the collection of student biometrics. Gilad has consulted on trust issues for the UK government's identity assurance program, Verify.gov, and is the author of *Privacy and the Internet of Things* (O'Reilly Media, 2017).

Gilad has a 20-year career in IT, having worked with identity technologies, digital media, automation, and telecommunications. Prior to becoming a researcher, he helped design, prototype and manufacture the world's only robotic video migration system, known as SAMMA, which won an Emmy Award for technical and engineering excellence in 2011. Gilad is a member of the UK Cabinet Office Privacy and Consumer Advisory Group, which provides independent analysis and guidance on Government digital initiatives, and a member of the IEEE Global Initiative for Ethical Considerations in Artificial Intelligence and Autonomous Systems. He is a Visiting Scholar at the UC Berkeley School of Information and a Visiting Researcher at the Horizon Digital Economy Research Institute. Gilad is a graduate of the University of Chicago, NYU and the University of Nottingham.

ERIN KENNEALLY, J.D.

Erin Kenneally is a currently a Program Manager in the Cyber Security Division within the U.S. Department of Homeland Security Science & Technology Directorate. Her portfolio comprises trusted data sharing and research infrastructure (IMPACT - Information Marketplace for Policy and Analysis of Cyber-risk and Trust), Data Privacy, cyber risk economics (CYRIE - Cyber Risk Economics), and information & communications technology ethics. Kenneally is Founder and CEO of Elchemy, Inc., and served as Technology-Law Specialist at the International Computer

Science Institute (ICSI), the Center for Internet Data Analysis (CAIDA) and Center for Evidence-based Security Research (CESR) at the University of California, San Diego.

Erin is a licensed Attorney specializing in strategy, research and execution of challenging and emergent IT legal risk solutions. This including translating, transitioning and implementing applied R&D and ops solutions at the crossroads of IT and advanced analytics, law, and ethics. She holds Juris Doctorate and Masters of Forensic Sciences degrees, and is a graduate of Syracuse University and The George Washington University.

ACKNOWLEDGEMENTS

We gratefully acknowledge the generous financial support for this project provided by The William and Flora Hewlett Foundation.

We are deeply grateful to the colleagues, friends and institutions that supported the creation of this work, and the forty experts, practitioners and researchers who graciously gave us their time to be interviewed and participate in our workshops. We would like to thank particularly Eli Sugarman, our Program Officer at the William and Flora Hewlett Foundation, and his colleagues for being willing to fund this report, and for providing an excellent venue for our first workshop. Also, we are grateful to Steve Weber, Betsy Cooper, Kristin Lin, Chuck Kapelke and Allison Davenport of the Center for Long-Term Cybersecurity at UC Berkeley for their encouragement, advice and administrative support, and for publishing a concise white paper based on this report. Thanks as well to Ian Wallace and New America for hosting our DC workshop in their beautiful offices. We'd like to thank the Advisory Board of the IoT Privacy Forum: Michelle Dennedy, Jim Adler, Sharon Anolik and Scott David. We're extremely grateful to our friends and colleagues who reviewed sections of this report for their time, encouragement and thoughtful critiques: Dr. Ewa Luger, Jo Breeze, Nik Snarey, Garreth Niblett, Steve Wilson, Dr. Ben Zevenbergen, Ian Skerrett, Steve Olshansky, Claire Milne, Dr. William Lehr, Dr. Tyler Moore, Shazade Jameson, Robin Wilton, Kasey Chappelle, Liz Coll, Emma Lindley, Dr. Chris Pounder, Sudha Jamthe, Josh Rubin, Dr. Richard Gomer, Rose Keen, Aoife Hagerty, Carlos Ruiz-Spohn, Elliotte Bowerman, Kaitlin Mara, Henrik Chulu, and Danielle Pacheco.

TABLE OF CONTENTS

INTRODUCTION

"Modern civilization is naturalistic, mechanistic, its rhythm the tempo of machines, each one of which is a creature of problem-solving intelligence. It is an unstable equilibrium of forces, the shifting patterns of which require of mankind ever more insight and calculation."

— Everett Dean Martin, 1928[1]

Scarcely a week goes by without a significant privacy event, data breach or legal decision commanding the attention of journalists and the public. Perhaps this is unsurprising: privacy debate has historically gone hand in hand with new technical developments, and given the rapidity of change all around us it makes sense that the stories are flying fast and furious. Given this increased rate of technological change and media interest, it is more crucial than ever to ensure depth and nuance in our privacy discourse. It is this belief in the value of nuanced discussion that has animated the research for this report about privacy and the Internet of Things (IoT).

There have been many names for the IoT over time: ubiquitous computing, ambient intelligence, machine-to-machine communications, pervasive computing, and, most recently, cyber-physical systems. The terms emerged from various disciplines, but they all point in the same direction. These persistent attempts to find a suitable term for the phenomenon reveal an awareness that the world is in rapid transition towards more comprehensive monitoring and connectivity, that this will likely have a profound impact on our lives, and that it is important to start anticipating the potential consequences. Our physical and informational world is evolving, and with it, the concept of privacy as we know it. The potential ramifications of the advances of the Internet of Things are a matter of profound concern to many people.[2]

1 Martin, 1928, p. 364

2 E.g., see the Pew Research Center's "The state of privacy in post-Snowden America: What we learned," available at https://pewrsr.ch/2HY6sps, and findings from the EU-funded CONSENT project, "What consumers think," available at http://bit.ly/EUConsent

This report is the culmination of eighteen months of research on how the Internet of Things affects and will affect privacy, and vice versa. It is exploratory in nature, attempting to test preexisting ideas and surface new ones. It presents new empirical data gathered through seventeen interviews and two workshops with a total of forty experts, scholars and practitioners, plus an extensive literature review. The gathered data was coded and analyzed through thematic content analysis techniques.[3] Throughout the report, we include many quotes from our interview and workshop participants to convey the richness of their views and voices.

Our intended audience for this report encompasses the industry, grant funding, technologist, privacy, policy, and academic communities. Writing for such a broad audience is a tremendous challenge, and we hope to have struck a balance between detail and brevity, flow and terminology. The title of this report reflects two aspects of the IoT that we highlight throughout the report: the transparency of people and places that the IoT engenders, and the opacity of the devices themselves and the data flows and organizations behind them.

This report is divided into the major themes that emerged from our interviews, workshops and literature. They are: *Boundaries*, *Governance*, *Privacy Management*, *Incentives*, *and Risks of Harm*. These sections are followed by an overview of *Emerging Frameworks and Strategies* to improve IoT privacy, and then our *Recommendations*, which augment the prior section. To facilitate further research and exploration of IoT privacy topics, the appendices contain lists of relevant academic departments and research institutes, conferences, and protocol and standards efforts.

3 Aberbach and Rockman, 2002; Braun and Clarke, 2006; Saldaña, 2009; Vaismoradi et al., 2016

DEFINITIONS

There is no strict definition of the Internet of Things. The term is a catch-all for the proliferation of objects in our homes and workplaces and cities that are acquiring varying degrees of networked intelligence. Devices that sense and communicate are not new, but technological developments have made sensing and connectivity inexpensive, unobtrusive and ubiquitous.

IoT refers to the increasing emergence of devices that are 'smarter' than they've historically been: a thermostat that knows a homeowner's preferred temperature, a watch that tracks fitness and can locate its wearer via GPS, a car that drives itself.

DEFINING THE INTERNET OF THINGS

The IoT is characterized by several converging trends: ubiquitous network access, inexpensive sensors, computation and storage, miniaturization, location positioning technology, and the advent of smartphones as a platform for device interfaces. In this way, the IoT should be seen as evolution in product development.

For the purposes of this paper, the IoT is defined as the collection of devices that have the ability to sense, amass, and analyze data and to communicate through networks. These devices might be found in the home, such as smart lighting or virtual assistants or TVs with cameras and microphones; outdoors in public, such as smart electricity grids, adaptive traffic signals and street lighting with gunshot detectors; in particular industries, such as remote monitors for health conditions or personally tailored advertising; in retail environments, detecting who and where people are in shops, observing where they look or linger; or on a user's person, such as wearable fitness trackers or head-mounted, networked cameras.

What unites them is that they are not full-fledged computers, but rather purpose-built devices, most of which are familiar, e.g., a vacuum cleaner or an ingestible pill. But these familiar objects have been upgraded with an ability to gather, process and transmit information, thus becoming new actors in the informational world.

Given the high diversity of connected devices, separating them into categories can assist in conceiving of and analyzing privacy implications. It should be noted that these categories are not absolute and overlap in many ways:

Ecosystem: This approach arranges the IoT by technical elements: Operating Systems, Platforms, Device Types, Infrastructure, Protocols, Processors, Interoperability, Standards[4]

Data Flow: The IoT here is viewed from the perspective of directional data flows. Each layer comes with its own stakeholders, commercial arrangements, and interoperability issues: Things >> Communications >> Gateway >> Data >> Integration >> Consuming Applications[5]

4 Cloud Security Alliance Mobile Working Group, 2015

5 Ovum, 2015

Technical Layers: The IoT can be separated into different layers from the perspective of security: network layer, application layer, device level, physical layer, human layer[6]

Industry Verticals: This describes the IoT according to markets for goods or services. Common arrangements include: Healthcare, Logistics, Energy & Utilities, Public Infrastructure & Services, Construction, Transportation, Retail, Media, Insurance, Entertainment, Telecommunications, Education, Banking, Law Enforcement, Agriculture, Consumer Goods

Contexts: This alignment considers the IoT according to its deployment settings: home, workplace, retail environments, personal vehicle, public transportation, in public, borders, government interactions, law enforcement interactions

The all-encompassing nature of the categories above shows that the Internet of Things will in the not-too-distant future merely be Things. That is, the IoT is today's label for the current step in the evolution of technology products. We believe the term IoT will, sooner rather than later, go away, much as 'mobile computing' did.

6 Cloud Security Alliance Mobile Working Group, 2015

DEFINING PRIVACY

It is clear that such devices will have an impact on privacy, but, as we shall show, there is debate as to whether they create new privacy issues or whether their effects on privacy merely mirror preexisting concerns. Exactly what is meant by privacy varies widely in different legal contexts, different cultures, and from person to person.

DEFINING PRIVACY

We find value in Alan Westin's classic definition of privacy as "the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others."[7] And, though it may be worn possibly to the point of being threadbare, Warren and Brandeis' conception of privacy as a 'right to be let alone'[8] is still useful to bear in mind, especially as they envisioned this right to encompass thoughts, emotions and sentiments, which is particularly germane to the IoT. We also find useful Westin's view that privacy protects four 'states': solitude, intimacy, anonymity, and reserve.[9] That said, these views are predicated in part on harms resulting

7 Westin, A., 1967, p. 7
8 Warren and Brandeis, 1890
9 Westin, 1967, p. 31

from *invasion*. We argue at different points in this report that the IoT threatens to decompose the notion of privacy invasion because of increasingly omnipresent sensors and because many IoT devices will be *invited* into our lives.

While the United States' legal and, arguably, cultural approach to privacy is largely organized around individualistic harms, other vital aspects of privacy appear both in the US and, to a greater degree, in Europe. The power and breadth of the right to privacy is one of the major differences between Europe and the US. Europe's somewhat more precautionary privacy laws hold that privacy and data protection are fundamental rights, whereas in the US a clear showing of harm after a privacy violation is the main trigger for legal intervention. Another essential dimension to consider is *expectation*, or how people expect that their information will be collected or used. In the US, the 'reasonable expectation of privacy' is a principal standard used to judge whether someone's privacy has been violated. The obvious problem with this standard is that expectations can change over time, and powerful actors can deliberately shift them. Nevertheless, one important privacy theory argues that when informational norms about the context in which information is disclosed or used are breached, as when data leaks from one context to another (e.g., from home to the workplace), people experience this unexpected boundary crossing as a privacy violation.[10] Related to this are theories of *boundary management*, which claim that privacy is part of people's ability and need to negotiate the boundaries between themselves and others, and with society at large. We discuss both the contextual and boundary dimensions of IoT privacy at length in the *Boundaries* section.

PRIVACY, THE SELF, AND DEMOCRACY

Broader yet is the idea that privacy is vital to the construction of the self and the health of a democracy. In 1998, the early days of the commercial internet, scholar Phil Agre wrote, "control over personal information is control over an aspect of the identity one projects to the world, and the right to privacy is the freedom from unreasonable constraints on the construction of one's own identity."[11] Classic formulations of liberal democracy venerate *autonomy* and *self-determination* as core aspects of what it means to be free. Agre connects how people's control over their *data impression* is linked to their ability to freely, autonomously construct themselves, in conjunction with the idea that privacy protects people's ability to make decisions free from unwanted interference. In other words, privacy

10 Nissenbaum, 2009
11 Agre and Rotenberg, 1997, p. 7

assists freedom of thought and of action. These freedoms, along with the ability to dissent, to protect deliberation from undue commercial or government influence, to speak freely and sometimes anonymously, and to think and behave in ways that may deviate from community norms are essential elements of participating in democracy. Consider the simplicity of voting in secret behind a curtain: privacy protects people from retribution for their political choices. It preserves the belief that those choices and the thoughts that led to them are their business and theirs alone.

Privacy is a *social* and *collective* value; a vital counterweight to earlier views that privacy is concerned with harms to and capabilities of *individuals*. If privacy protects people's capacities to participate in democracy, then it confers benefits on society as a whole. Scholar Priscilla Regan argued that privacy is a *common* value: "although different people exercise the right to free conscience differently, believe in different things, and belong to different religions, all individuals have a common interest in this right. The same is arguably true for privacy."[12] The preservation of privacy must therefore be enacted at social levels and not left exclusively to the domain of individual people and how they experience it. It is "hard for any one person to have privacy without all persons having a similar minimum level of privacy."[13] Other scholars argue that privacy is *constitutive* of society,[14] integrally tied to its health, and that privacy is a public good.[15] In this way, privacy regimes can be seen as social policy, encouraging beneficial societal qualities, discouraging harmful ones, and safeguarding democracy.[16] Market actors like device manufacturers and service providers are essential contributors to how privacy manifests, both in the sense of mechanisms (settings, controls, switches) and norms (what is and is not acceptable practice, when to show notices, sharing defaults). But these manifestations occur largely according to commercial logic; business is concerned with sales. The long-term health of democratic society, embodied in the preservation of social values like fairness, freedom of thought, and protection of the vulnerable, is decided in political and policy realms.

The way we discuss privacy, the way we employ it to govern information and the power it holds, and the way we encode it into formal and informal policy instruments have direct bearing on the kind of society we collectively create.

12 Regan, 1995, p. 221

13 Regan, 1995, p. 213
14 Schwartz, 1999
15 Fairfield and Engel, 2015
16 Bennett and Raab, 2003, Part 1

BOUNDARIES

One of the main themes to emerge from our interviews, workshops, and literature is that **connected devices challenge, cross and destabilize *boundaries***. These can be physical boundaries, such as the walls of the home or the skin; data-type boundaries, such as personally-identifiable versus non-identifiable; or regulatory boundaries, such as telecommunications, FTC jurisdiction, or airspace regulation. Related to this is **the collapsing of *social contexts***: family life, work life, the inner life, health care, education, offline vs. online, public vs. intimate.

This section explores how the IoT affects boundaries and people's ability to manage them; especially as we are often *inviting* these devices into private spaces rather than being *invaded* by them. Further, the scale and proximity of sensors being deployed in the human environment will begin to challenge historical notions of bodily and emotional privacy, especially given an increase in commercial interest in emotion data. We discuss the boundaries of the home and the body, sensor fusion, and the commercial drive to cross contexts. We close by presenting two theories raised by our participants that are pertinent to how people perceive and regulate privacy: contextual integrity and boundary management.

THE HOME

Our interview and workshop participants observed that emerging connected devices challenge the physical and contextual boundaries of the home.

Commenting on the impact of the IoT on the perceived sanctity of the home, Heather Patterson of Intel remarked, "I've been thinking lately about how the IoT has the potential to really shift… the home from a black box, what used to be a protective, safe space, to more of a glass house where everything that we do is now readily apparent to people who are willing to look for it."[17] One security researcher agreed that the IoT is redefining the boundaries of the home, noting:

Our mental boundaries in our homes and private lives change with the presence of IoT. We have a legal framework based around our home, geographic boundaries on data territoriality, but the IoT traverses these boundaries.[18]

A key technology crossing these boundaries is the smart speaker/virtual assistant: Amazon's Echo with Alexa, Microsoft's Cortana, Google Home, and the recently released Apple Home Pod with Siri. These devices introduce a combination of microphones, artificial intelligence, voice recognition, and the melding of personal profile information gleaned from the use of other services.

17 Interview
18 Workshop comment

These devices are qualitatively different than, say, televisions with voice recognition because of the degree of AI combined with the combination of extensive profile information. There's no doubt this new class of technologies brings pleasure, convenience, entertainment and a platform for new voice interactive applications. Rather, the issues worth exploring relate to the placement of a general-purpose microphone – and increasingly cameras as well – into the home, a context classically seen as the quintessential private space.

Privacy in the home is an embodiment of *privacy of location*, "the right of an individual to be present in a space without being tracked or monitored or without anyone knowing where he or she is."[19] The home also embodies *spatial privacy*: "the protection of the privacy of people in relation to the places where they enact their private life. Classically, this is the dwelling or house, but it can stretch to other 'places of private life' … private places with discernable boundaries."[20]

Regarding the governance of privacy in the home, Justin Brookman, formerly of the FTC's Office of Technology and Investigation and now Director for Privacy and Technology at the Consumers Union, noted, "Privacy in the home is a very bipartisan thing. Think about Scalia: the man's home is his castle."[21] Brookman is referring to Justice Antonin Scalia's opinion in the seminal US Supreme Court case, *Kyllo v US*, where police used a thermal imager without a warrant to detect heat patterns emanating from the defendant's house from marijuana cultivation. Justice Scalia opined:

> *In the home… all details are intimate details, because the entire area is held safe from prying government eyes… We have said that the Fourth Amendment draws "a firm line at the entrance to the house."*[22]

The 2001 case specifically concerns law enforcement´s ability to breach the home boundary, but it is illustrative of the legal and cultural sanctity of the walls of the home. In Europe, this sanctity is embodied in Article 7 of the Charter of Fundamental Rights of the EU: "Everyone has the right to respect for his or her private and family life, home and communications."[23] Two US states, Arizona and Washington, guarantee privacy in their constitutions,

19 Wright and Raab, 2014, p. 7
20 Koops et al., 2016 p. 28
21 Interview
22 Kyllo v. U.S., 533 U.S. 27 (2001)
23 European Union, 2012, Art. 7

both citing the home context: "No person shall be disturbed in his private affairs, or his home invaded, without authority of law."[24] Further, psychological, sociological, anthropological and other literatures have argued that the home is a foundational aspect of the human experience; "a physical space that is lived – a space that is an 'expression of social meanings and identities.'"[25]

Much of the legal focus of privacy of the home centers around *invasion*, and arguably so do the location and spatial privacy conceptions above. But the introduction of virtual assistants, networked toys and interactive televisions is based on *invitation* – people willingly purchasing and installing these products at home. To describe some of the new challenges posed by inviting IoT devices into private spaces, Bastian Könings, Florian Schaub and their colleagues have been expanding upon 'territorial privacy,' a concept that overlaps with both spatial and location privacy:

> *Whereas in a traditional scenario the physical boundaries of a room would also mark the boundaries of a user's private territory, this situation will drastically change in future ubiquitous environments. Invisible embedded sensors, actuators and in particular wireless communications could widen the boundaries of a private territory far beyond its physical boundaries. As a consequence* ***the ability to perceive and control who is observing or disturbing a user in her private territory will decrease or even cease to exist.***[26]

This view of privacy helpfully moves discussion from intrusion to control. The authors write: "The goal of territorial privacy is to control all physical or virtual entities which are present in the user's *virtual extended territory* in order to mitigate undesired *observations* and *disturbances*, and to exclude undesired entities from the *private territory*."[27]

Given the likelihood of ubiquitous data collection throughout the human environment in the near future, **the notion of invasion may decompose**; all the more so as people's expectation of being monitored increases. It is therefore crucial to continue to introduce privacy approaches that assert control over devices and data flows. The *Emerging Frameworks and Strategies* section explores practical efforts to enhance user awareness and control.

24 AZ Const., Sec. 8; WA Const., Sec. 7
25 Mallet, 2004, p. 80, citing Wardaugh
26 Könings and Schaub, 2011, p. 105, emph. added
27 Könings et al., 2016: 136, orig. emph.

THE BODY AND EMOTIONAL LIFE

IoT devices are slowly beginning to breach another sensitive boundary: the skin, in the sense of the boundary between people's inner and outer lives – the domain of *bodily privacy*. As with the home, legal and cultural sensitivities exist around the sanctity of the body and mind. Consider prohibitions on forced medical procedures or the surrendering of blood samples[28] and restrictions on administering lie detector tests.[29] The Internet of Things will likely start to test the sensitivities around bodily privacy and the revered private nature of one's thoughts.

The IoT's increased monitoring of human activity is fueled by *scale* – a greater number of sensing devices and sensor types – and also by the *proximity* of those devices to people's bodies. The examples are obvious: mobile phones with their cameras and microphones are constantly within arm's reach; fitness trackers have diversifying biometric sensors; Nest has expanded from networked thermostats to indoor surveillance cameras;[30] toys are listening to children.[31] The trend is clear: the commercial market is offering ever more devices to monitor people's activities, environments, and, importantly, their physical bodies and emotions.

Cameras and microphones are general-purpose sensors, and the platforms they connect to continue to expand their capabilities. Computer vision can noninvasively detect facial blood flow and other biological artifacts.[32] Wearables go a step further, employing a variety of both biometric and non-biometric sensors like GPS, accelerometers, altimeters, and gyroscopes. Sensor component costs continue to fall so that technology once constrained to medical and industrial settings is now proliferating in consumer

28 Sarnacki, 1984
29 Regan, 1995, p. 144-169

30 See https://nest.com/cameras/
31 See, Hello Barbie http://hellobarbiefaq.mattel.com/
32 Sikdar et al., 2016

contexts. **These technologies and trends have contributed to a rise in emotion detection and 'affective computing'** – computing that relates to, arises from, or influences emotions.[33]

Emotion detection is the capture and analysis of facial data (surface and below), biometrics and voice, plus sentiment analysis which is drawn from textual and graphic expressions. Prof. Andrew McStay, author of the book *Emotional AI: The Rise of Empathic Media*, explained that in the wearables domain, sensors can measure everything from respiration to blood flow, skin conductivity to EEG rates. The data patterns generated by these measurements are then indexed according to the emotion they ostensibly represent. This presupposes that the collected data can be mapped reliably to emotions, but McStay notes that such claims should be viewed cautiously. Despite the scientific research and findings touted by these companies, neuroscientific, social psychology and humanities literatures have yet to come to a consensus as to what emotions actually *are*. Nonetheless, development of and commercial investment in emotion detection technology continues unabated:

We're seeing a net-rise of interest in sentiment and emotion capture. *The industries are wide ranging: automobiles, insurance, health, recruitment, media… basically anywhere it's useful to understand emotions. In terms of industries that are taking the lead on this, advertising and marketing are the obvious ones. Increasingly we're seeing retail move in to that area as well, plus all sorts of different sectors, ranging literally from sex toys all the way up to national security agencies, and all the marketing and organizational stuff in-between.*[34]

In late 2017, a consumer advocacy group published research on a range of patents secured by Google and Amazon relating to potential future functions of their digital home assistant products.[35] In one of these, Amazon patented a method for extracting keywords from ambient speech which would then trigger targeted advertising. The method is to listen for specific kinds of keywords – verbs expressing interest or desire: "like," "love," "prefer," "hate," "enjoy," and so on.[36] A person may say, "I love skiing," and then be served relevant ads, even though the speech was spoken to another

33 Picard, 1995, p. 1

34 McStay, Interview, emph. added
35 Consumer Watchdog, 2017; see also Maheshwari, 2018
36 Edara, 2014

person, not the virtual assistant. In another patent, Google describes a smart home in which "mischief may be inferred based upon observable activities" by children.[37]

It seems highly likely that companies will continue to expand into emotion and sentiment observation, gaining ever more access to what lies below our public behavior and speech. We further explore McStay's link between marketing, emotions and consumerism in the *Risks of Harm* section. For the moment, it suffices to say that **emotion capture raises questions of bodily privacy.**

Privacy scholar Gary T. Marx argued that a privacy intrusion occurs when the physical barrier of the skin is crossed:

> *Informational privacy encompasses physical privacy. The latter can refer to insulation resulting from natural conditions such as walls, darkness, distance, skin, clothes and facial expression. These can block or limit outputs and inputs. Bodily privacy is one form of this. This is seen in crossing the borders of the body to implant something such as a chip or birth control device or to take something from it such as tissue, fluid or a bullet.*[38]

However, professor of philosophy Judith DeCew, referring to Warren & Brandeis' landmark "Right to Privacy" article,[39] points out that even as early as 1890 a less corporeal, less physically transgressive reading of the bodily privacy right was appropriate. She noted that Warren and Brandeis had proposed a "more general right to privacy which would protect the extent to which one's thoughts, sentiments, and emotions could be shared with others."[40]

Until recently, the boundaries of the skin and the inner life have arguably been less of an information privacy concern.[41] However, the scale and proximity of IoT sensors is reinvigorating this 130-year-old privacy interest. While we believe that the IoT mainly represents an amplification of existing data collection trends, large-scale emotion detection is a significant shift. We expand on the particular risks of emotion detection in the *Risks of Harm* section.

37 Fadell et al., 2014
38 Marx, 2012, p. x
39 Warren and Brandeis, 1890
40 DeCew, 2013
41 An exception to this is the privacy concerns of biometrics collection

SENSOR FUSION

The combination of multiple kinds of sensor data into a more revealing picture has been termed 'sensor fusion.' We view sensor fusion as a weakening of boundaries between data types and data sets; the amalgamation of data from different contexts into a more revelatory picture. Prof. Scott Peppet of the University of Colorado Law School writes:

Just as two eyes generate depth of field that neither eye alone can perceive, two Internet of Things sensors may reveal unexpected inferences. *For example, a fitness monitor's separate measurements of heart rate and respiration can in combination reveal not only a user's exercise routine, but also cocaine, heroin, tobacco, and alcohol use, each of which produces unique biometric signatures... [E]ach type of consumer sensor... can be used for many purposes beyond that particular sensor's original use or context, particularly in combination with data from other Internet of Things devices.*[42]

The quality and quantity of IoT sensor data enhances the inferences that can be drawn from that data, while also increasing the chances that people can be identified from such data, even when it's analyzed in aggregate.[43] Lee Tien, Senior Staff Attorney at the Electronic Frontier Foundation, observed that not only is the fusion of data from different devices poorly disclosed by companies, it is also poorly understood by people.[44] As such, the use of connected devices can lead to unexpected revelations to the myriad companies in the data supply chain of IoT information and their business partners. We turn now to the impacts of violated expectations.

42 Peppet, 2014, emph. added, references omitted

43 Kohnstamm and Madhub, 2014; Peppet, 2014, p. 129-131

44 Interview

THE IMPACT OF DESTABILIZED BOUNDARIES

There are important privacy implications when devices cross boundaries in unexpected ways. In particular, this emerging characteristic of IoT devices challenges *contextual integrity* and *boundary management*. Contextual integrity is a theoretical framework that tries to explain why people feel privacy violations.[45] It describes a set of informational norms that govern the varying contexts of social life: norms concerning the transmission or distribution of personal data, the type of data, and the actors involved. The framework explores the way we manage different social contexts, such as home life, work life, medical care, education and commerce:

[P]eople act and transact in society not simply as individuals in an undifferentiated social world, but as individuals in certain capacities (roles), in distinctive social contexts, such as health care, education, employment, the marketplace, and so on. These contexts should be understood as structured settings whose features have evolved over time — sometimes long periods of time — subject to a host of contingencies of place, culture, historical events, and more.[46]

When the norms governing information exchange within particular contexts are violated or contextual boundaries are unexpectedly crossed, people experience it as a privacy violation. For example, a person expects her doctor to disclose information to other medical staff and her insurance company, but not her family; a person expects that family members in the home know which television programs she watches, but does not expect her boss to know. In her 2009 book, *Privacy in Context*, sociolegal scholar Helen Nissenbaum summarized, "What people care most about is not simply *restricting* the flow of information but ensuring that it flows *appropriately*..."[47]

45 Barth et al., 2006; Bruening and Patterson, 2016 ; Nissenbaum, 2009

46 Barth et al., 2006, p. 184

47 Nissenbaum, 2009, p. 2

Contextual integrity concerns itself with the boundaries of social contexts. In this way, it bears kinship to the theories of 'boundary management.' In 1976, social psychologist Irving Altman conceived of privacy as "an interpersonal boundary control process" and "selective control of access to the self or one's group."[48] In contrast to traditional ideas of privacy as the act of retreating or hiding information, Altman called privacy a "dynamic dialectical process":

> ***Privacy is a boundary control process*** *whereby people sometimes make themselves open and accessible to others and sometimes close themselves off from others...* ***[T]he issue centers around the ability of a person or a group to satisfactorily regulate contact with others.***[49]

Importantly, Altman distinguishes his privacy framework from others that rely primarily upon withdrawal, solitude, anonymity and secrecy, arguing instead that *selective sharing* of the self co-occurs with withdrawal and reserve. In this sense, Altman's theory of boundary control ties into Nissenbaum's assertion that people are more concerned with controlling what gets shared than by eliminating the sharing process altogether. Altman summarized, "Privacy mechanisms serve to define the limits and boundaries of the self. By being able to change the permeability of the boundaries around oneself, a sense of individuality develops – sometimes incorporating others and the world, sometimes keeping them out."[50] In this, Altman argues that privacy serves the goal of promoting autonomy.[51]

Later privacy and human-computer interaction (HCI) scholars incorporated and expanded on Altman's views. A seminal HCI paper by Palen and Dourish observed, **"Privacy management is not about setting rules and enforcing them; rather, it is the continual management of boundaries between different spheres of action and degrees of disclosure within those spheres."**[52] By blurring boundaries – between contexts, self and other, the home and public, expected islands of reserve and visibility to third parties – IoT devices challenge people's ability to negotiate those boundaries.

Some of our respondents identified contextual integrity and boundary management as issues particularly problematic to IoT:

48 Altman, 1976, p. 7-8
49 Altman, 1977, p. 67-68, emph. added
50 Altman, 1976, p. 26
51 Similar to Agre; see the Definitions section
52 Palen and Dourish, 2003, p. 131

If you have sensors integrated into buildings, into homes, if you have smart devices like smart speakers that are basically listening to you all day long, ***it becomes really difficult to know what information's being collected, when it is being collected, and it's very opaque where that information flows and which entities might get access to [it]****... It's no longer sufficient to close a door if you want to have a private conversation – there might be other sensors that are picking up who's in the room, what they are talking about, what is the mood of these people.*[53]

You're putting your smart lights into your office building, but that also lets you track employee productivity and usefulness. Creating boundaries between a productive work environment and an incredibly hyper surveillance work environment is really tough, and probably only getting worse.[54]

Boundary work is managing those different aspects of life, of tasks, of kids and work and environments and social context... When we're surrounded by other peoples' stuff that's collecting who knows what, it makes that boundary work really challenging.[55]

Similar concerns arise in recent published work. Margot Kaminski cites the contextual integrity and boundary management issues with household robots, raising concerns about people suppressing their own speech and conforming their behavior:

When information revealed in the home is shared and used outside of the home, people may stop trusting that the home is a private location, and may stop sharing information and conform their behavior to majority norms even within the home.[56] *... Household robots threaten the ability of individuals to conduct... 'boundary management' because in addition to crossing physical boundaries, or being able to 'sense' through physical boundaries... robots' social features may elicit trust where trust is not deserved...*[57]

Challenges to boundary management, however, affect not only personal privacy, but society at large. The European Commission's IoT Expert Group worried about how the breakdown of contexts could affect democracy:

53 Schaub, Interview, emph. added
54 Jerome, Interview
55 Jones, Interview

56 Kaminski, 2015, p. 664, emph. added
57 Ibid.

The perimeter of a context, keeping certain information or actions restricted to the boundaries of a particular restricted type of interaction, may silently disappear by technology that is as ubiquitous and interconnective as IoT. Such de-perimeterisation associated with converging technologies challenges the checks and balances associated with the separation of powers in our democracy.[58]

Prof. Julie Cohen cited the necessity for functional boundary management in order for people to determine the shape and course of their lives:

Privacy shelters dynamic, emergent subjectivity from the efforts of commercial and government actors to render individuals and communities fixed, transparent, and predictable. It protects the situated practices of boundary management through which the capacity for self-determination develops.[59]

Law scholar Anne Uteck related the degradation of boundary management to the breakdown of the reasonable expectation of privacy standard, and the threat of amplified, unwanted exposure:

Embedded surveillance networks undermine the pretense that we control our environment or our boundaries within it, a pretense fundamental to the traditional construct of privacy. *It is no longer just a question of further eroding the distinction between private and public spaces, or even simply blurring the boundaries of this distinction, but* ***the extent to which [ubiquitous] technologies have almost limitless potential to contravene any reasonable expectations of privacy in private and in public.***[60]

58 van den Hoven, 2016, p. 13
59 Cohen, 2013, p. 1905

60 Uteck, 2013, p. 82-83, emph. added; see also Reidenberg, 2014, p. 145

Contextual integrity and boundary management are two useful frameworks for exploring both the personal and the social implications of the context-crossing qualities of IoT devices. People's ability to negotiate the boundaries of their social and physical contexts affects their ability to manage their privacy. As that ability weakens, people may trust objects when they should not, and lose more of the ability to know who is gathering data about them. At the societal level, the porousness of context boundaries can have an effect on autonomy and democracy itself. How people behave is influenced by their ability to negotiate boundaries and know who is 'inside' and who is not.

CONCLUSION

There is undoubted value in digital devices and services 'following' people as they move through the different social spheres of their lives, predicting behaviors and offering suggestions and timely information. Indeed, one could argue that this is the 'natural' direction of personalizing user experiences. However, there are signs that the home is gradually losing its sanctity as a place of reserve because of the monitoring devices people are inviting in. Similarly, while content tailored to people's emotional states is attractive for gaming, entertainment and other uses, sensor scale and proximity fuels concerns over emotional and bodily privacy. Adding to this is sensor fusion, which, by combining data from disparate collections, can cause greater, and possibly unexpected revelation of personal traits and activities. The boundary-spanning quality of IoT devices has a privacy impact on both individuals and society at large.

GOVERNANCE

Governance of the IoT is intertwined with questions about how privacy rights and risks are managed by public, private and social institutions. The IoT commingles new and existing components and platforms, so unsurprisingly there is debate over whether it is an evolutionary or revolutionary development. As such, the extent to which the IoT warrants new laws, policies, and regulations is a key question.

Addressing privacy regulations for connected devices surfaces discussions about: 1) whether privacy issues are different in degree or kind from those covered by our existing Internet regulations; 2) how the IoT challenges the boundaries and enforcement of existing governance instruments; 3) political debates pertinent to the relationship between market innovation and government regulation; and 4) the level of maturity of social discourse about the interrelationship of technology and privacy.

OLD WINE IN NEW BOTTLES?

Is the Internet of Things different from the internet? Is there a sufficient difference to warrant changes in existing policies that govern informational privacy and data protection? The answers to these questions are pivotal in considering how to govern the IoT. Our research participants, literature, and stakeholders in the public policy process have a range of responses.

On one hand are those who consider the IoT a different beast. Joris van Hoboken of the Free University of Brussels, for example, sees IoT policy considerations in a broad sense, encompassing not just physical devices but society as a whole.[61] This view echoes that of Dr. Gerald Santucci, an expert in IoT and related policy, who observed, "The IoT does not concern objects only; it is about the relations between the everyday objects surrounding humans and humans themselves."[62] Florian Schaub of the University of Michigan noted the importance of the IoT's ability to affect the physical environment: "It's not just about information being collected about specific individuals, but it also becomes the question of technology becoming active in the user's environment - actuation."[63] A 2014 report by the Aspen Institute also highlighted how data sensors and autonomous actuators give the IoT unprecedented power in the physical, versus purely online, world:

61 Interview
62 Santucci, 2011, p. 84
63 Interview

To some extent, the IoT merely enlarges – vastly – the existing Internet. It dramatically increases the number of inputs into the systems of data collection and analytics that drive the digital world. In other ways, the IoT is something entirely new. It empowers our physical world to make decisions about how it interacts with us, without any direct human intervention.[64]

Prof. L. Jean Camp and her colleagues cited a similar view:

The challenges and opportunities arising from the IoT are fundamentally new, owing to the unprecedented combination of ubiquity, diversity, and connectivity among IoT devices, and ***the ability for many IoT devices to observe and actuate real world events without any explicit human interaction.***[65]

One security researcher also noted the issue of ubiquity and a lack of direct user interaction, seeing the IoT as more ubiquitous than mobile phones and embedded in the environment: "We wear them on our bodies, put them on our toys, they collect data in novel ways we didn't envision before. It's more of a passive question of sensors collecting information without our interaction with them." This is not to suggest that a regulatory response is required, he explained: "You could imagine, for example, proactive steps like an opt-in consent when a mobile device collects location information. That's not a regulatory requirement, but a self-regulatory norm that companies have adopted due to recognition of the potential harms."[66]

In 2017, the US Department of Commerce released its report, "Fostering the Advancement of the Internet of Things," gathering comments from industry, civil society organizations, advocates, trade bodies, academia and other regulatory agencies. Regarding privacy, it stated, "Several commenters argued that there are no new privacy issues related to IoT, that it is too early to craft regulatory responses, or that current regulation is sufficient."[67]

64 Goodman, 2015, p. 6-7
65 Camp et al., 2016, emph. added
66 Interview
67 US Dep't of Commerce, 2017, p. 31

On the subject of the necessity of crafting new regulations specific to the IoT, one senior government official responded:

I feel like every certain number of years, there's a new technology, whether it's the internet or mobile or IoT – we all of a sudden say, "Oh, we need to rethink everything in this context." I think that's just the evolution of technology.[68]

Ian Glazer, Vice President for Identity Product Management at Salesforce, also opined about the need for new policies:

*People tend to lose their ever-loving mind when they start talking about IoT. Everyone chill out. We have practice and principles that should be guiding us consistently, because if you have a good principle, it doesn't need to be changed for different modalities of technology. I get really nervous when people say, "Oh, we need a new **X** for this expression of technology." Well, if we had fundamental good practice, good principle, and good regulation, then it shouldn't have to get updated for a specific kind of technology.*[69]

Given this diversity of views about the novelty of the IoT, is it possible to reach an agreement as to whether it requires new kinds of policy instruments to govern it? Radical reconsiderations of policy can occur but are uncommon. Public policy theory and related areas of scholarship often argue that incrementalism is a dominant mode of policymaking.[70] Any realistic consideration of policy development must be seen in the larger context of existing policies, institutional arrangements, legislative appetite, and the politics of the day, which we discuss in the next section. These prior factors often give rise to 'path dependency' – a stability and sometimes irreversibility in the current state of policy choices.

We find that although this is largely the case with IoT policy, there are some instances of new policy choices ('path creation') within individual technology domains, such as the automotive sector. For example, a Senate version of a federal autonomous vehicle policy bill currently being debated in the US Congress calls for the creation of an online, searchable 'motor vehicle privacy database' which

68 Interview
69 Interview
70 Lindblom, 1979; North, 1991; Scott, 2003

would list all the personal data that vehicles collect, how that data and the conclusions drawn from it are used and disclosed, how long they will be retained, and when they will be destroyed.[71] In particular, the inclusion of language about disclosing the use of conclusions (read: inferences) drawn from interactions with vehicles is highly unusual in US information policy, and can be seen as a promising new path.

With regard to actuation, the legal and insurance communities are slowly beginning to grapple with questions of IoT liability. Given the recent first death from a self-driving car[72] and the world's largest distributed denial-of-service attack coming from compromised IoT devices in 2016,[73] issues of product liability and insurance are on the cutting-edge of IoT governance.[74]

71 S. 1885, 2017
72 Economist, 2018
73 Woolf, 2016
74 See Dean, 2018

GOVERNANCE AND INNOVATION PHILOSOPHIES

Public, commercial and policy discourse surrounding the IoT is suffused with talk of innovation; the need to encourage it, to prevent it from being stifled, to ensure society and its posterity benefit from it:

The Internet of Things has already led to important technological breakthroughs, and as it expands its reach, it has the potential to spur tremendous innovation. Our challenge is to find the proper balance between promoting this innovation and ensuring that our security and our privacy are protected as this valuable technology continues to grow.[75]

[B]ecause the Internet of Things... offers so many important economic and social benefits, countries should develop national strategies to promote its adoption and use.[76]

75 Internet of Things, 114th Cong., p. 4, 2015 (Introduction by Issa)
76 Castro and New, 2016, p. 1

One strand of discourse that encourages innovation while minimizing factors that could stifle it is dubbed 'permissionless innovation.' Adam Thierer of George Mason University, a central proponent of this idea, defines the idea variably:

The notion that experimentation with new technologies and business models should generally be permitted by default.[77]

The general freedom to experiment and learn through ongoing trial-and-error experimentation.[78]

Permissionless innovation is about the creativity of the human mind to run wild in its inherent curiosity and inventiveness. In other words, permissionless innovation is about freedom.[79]

Thierer contrasts this with (his version of) the 'precautionary principle,' which he criticizes for its constraining effect on innovation:

The precautionary principle holds that since every technology and technological advance could pose some theoretical danger or risk, public policies should prevent people from using innovations until their developers can prove that they won't cause any harms.[80]

Many critics adopt a mindset that views the future as something that is to be feared, avoided, or at least carefully planned. This is known as the 'stasis mentality' and it is, at its root, what motivates precautionary principle-based thinking and policymaking.[81]

77 Thierer, 2016, p. 1
78 Ibid., p. 8
79 Ibid., p. 9

80 Thierer, 2013, p. 353
81 Thierer, 2016, p. 23

This is a politically charged, narrow reading of the precautionary principle, portraying it as a force for prohibition. However, in a review of the precautionary principle in environmental regulation, legal scholar Richard Stewart laid out four categories of the precautionary principle:

"PP1. Scientific uncertainty should not automatically preclude regulation of activities that pose a potential risk of significant harm (Non-Preclusion PP).

PP2. Regulatory controls should incorporate a margin of safety; activities should be limited below the level at which no adverse effect has been observed or predicted (Margin of Safety PP).

PP3. Activities that present an uncertain potential for significant harm should be subject to best technology available requirements to minimize the risk of harm unless the proponent of the activity shows that they present no appreciable risk of harm (BAT PP).

PP4. Activities that present an uncertain potential for significant harm should be prohibited unless the proponent of the activity shows that it presents no appreciable risk of harm (Prohibitory PP)."[82]

Thierer is clearly aligning his view with the fourth category, protesting against the idea that developers should be required to prove a technology will do no harm before it is allowed to flourish. His response to this is inversion – arguing that regulation should not proceed without evidence of a potential harm:

The burden of proof should be on those who favor preemptive, precautionary controls to explain why ongoing trial-and-error experimentation with new technologies or business models must be disallowed.[83]

82 Stewart, 2002, p. 76; see also Sunstein, 2003
83 Thierer, 2016, p. 4

Industry positions on IoT governance are replete with this viewpoint:

US Chamber of Commerce: "Without evidence of heightened privacy concerns or consumer harm, there is no reason not to allow the IoT market to mature under the frameworks that exist for protecting consumers' legitimate privacy interests."[84]

T-Mobile: "NTIA and other federal agencies should only adopt cybersecurity and privacy regulations based on clear evidence of actual harm after balancing the need for regulation against the deterrent effects regulation will have on innovation and investment."[85]

Consumer Technology Association: "As a general matter, the increasing number of devices should not automatically trigger new regulations – before acting, there should be evidence of real harms."[86]

The very term 'permissionless' creates an antagonist; the government and concerned parties as bête noire to the industry and the market, whose noble goal to produce social good and consumer welfare through innovation must be defended. **The discourse of permissionless innovation fuels a general distrust in regulation by falsely pitting 'freedom' against reasonable efforts to govern technology and protect current and future citizens from economic, social and democratic harms.** However, as illustrated by Stewart, precautionary approaches exist that would allow for innovation while still mandating strong IoT privacy characteristics. Without a doubt, governance and legislation are imperfect mechanisms, and attempts to craft forward-facing consumer protection and privacy rules will be inevitably flawed. However, "the precautionary principle benefits privacy protection insofar as it emphasizes the normative values of prudence and transparency."[87] Prudence is warranted in particular when concerned with democratic harms because they are diffuse and take a long time to manifest and detect.

84 Eversole et al., 2016
85 Massey et al., 2016
86 Kearney and Reynolds, 2016

87 Costa, 2016, p. 23

Despite what Thierer would have us believe, outright champions of the prohibitory version of the precautionary principle are rare. Only a small number of technologies are regulated in this fashion, the most obvious being the development of new drugs and medical devices, and in transportation. Setting aside questions of regulatory efficiency, the reason for such prohibitory precaution is the potential threat to human life. But, to imagine that such prohibition exists commonly in other domains of technology regulation – in particular, consumer devices and internet technologies – is a straw man argument. In both the United States and Europe, there are almost no examples of consumer technology regulation employing a prohibitory version of the precautionary principle.

Luiz Costa of the University of Namur sees the precautionary principle as useful to privacy and data protection law because it addresses issues of information and power asymmetry, spurs public debate, recognizes that monetary damages for future harms may be inadequate, and takes account of the uncertainty of scientific evaluation.[88] On the first two points, Costa notes that legislation protects citizens by counterbalancing the power of governments and industry.

88 Costa, 2016

In part this is achieved through applying the precautionary principle **"in order to avoid risk-taking without a larger public discussion.** The promotion of this principle is therefore a way to involve citizens in decision-making."[89]

Costa cites the United Nations Rio Declaration on Environment and Development, which employs the precautionary principle to protect the environment: "Where there are threats of serious or irreversible damage, lack of full scientific certainty shall not be used as a reason for postponing cost-effective measures to prevent environmental degradation."[90] This formulation aligns with Stewart's PP1 and PP2 above. As to the inadequacy of monetary damages, Costa writes that:

> ***Precaution... considers that some damages cannot be repaired or compensated with money because not everything can be converted into money.*** *Considerable oil leaks can cause damages to the environment that are irreparable; lives lost in an accident are irreversible. Instead of compensating damage, precaution urges the need to avoid some damages.*[91]

89 Ibid. p. 15, emph. added
90 UN General Assembly, 1992
91 Costa, 2016, p. 17, emph. added

This is particularly salient to privacy and data protection, where the potential dangers to individuals are so often non-economic: embarrassment, stigmatization, loss of dignity, intrusion, and decisional interference.[92] This is also true for the longer-term, difficult-to-detect threats to autonomy, freedom of thought, private spaces, and a liberal democratic order. **As with the environment, some injuries to the health of a democracy are potentially irreparable.** Or, as Paul Ohm writes: "If we worry about the entire population being dragged irreversibly to the brink of harm, we must regulate in advance because hoping to regulate after the fact is the same as not regulating at all."[93]

Taken together, these arguments form a counter-narrative to the proponents of permissionless innovation, who are largely silent on threats to the health of a democracy, or the danger of a general diminishment of private spaces. We argue that the discourse of permissionless innovation is ultimately a deregulatory thrust, often libertarian in character, that serves the interests of economic actors while denying society the vital role of government regulation in protecting a broader range of social and democratic interests.

92 See generally Solove, 2006
93 Ohm, 2010, p. 1751

BLURRED REGULATORY BOUNDARIES

Overlapping the thematic areas of Boundaries and Governance is the view that the Internet of Things muddles the boundaries between regulatory agencies. Further, the US's current sectoral approach is seen by some stakeholders as being insufficient to deal with the many gray areas that arise in IoT governance policy.

There are a number of examples of regulatory blurring. Drones (of the non-hobbyist type) in the US are regulated by the Federal Aviation Administration, but they do not have the authority to regulate their privacy aspects. The National Highway Traffic Safety Administration (NHTSA) released guidance on automated vehicles in September 2016 with privacy guidelines[94] that included a recommendation for data minimization, but one year later removed all references to it in an update, saying instead that "the FTC is the chief Federal Agency charged with protecting consumers' privacy and personal information."[95]

The health sector is emblematic of muddled regulatory boundaries. Michelle De Mooy of the Center for Democracy & Technology (CDT) described its regulatory confusion:

If you've created a mobile health app that is spitting out outcomes based on the data you're generating, what does that mean? Where are the lines for what should be regulated? Either because it's creating health recommendations that would affect your health and safety, or because it's promising things that it can't deliver – that would be the difference between the Food & Drug Administration or the FTC being involved. The Department of Health and Human Services, the FTC, the FDA have all recognized that this is rising: this blurring of the uses for sensors and IoT and health-related care are blending over, at least in a regulation sense.[96]

Lee Tien of the Electronic Frontier Foundation sees US federal policy being hampered by agency competition, a lack of regulatory cohesion, overlapping regulatory jurisdictions, and a lack of familiarity with emerging technology: "There is a lot of jockeying… that's exacerbated by the fact that a lot of these issues are new, so they haven't had to deal with privacy before."[97] Relatedly, the US Department of Health and Human Services noted how both mobile health developers and the citizenry could be misled by the absence or overlap of regulatory authorities:

94 NHTSA, 2016
95 NHTSA, n.d.
96 Interview
97 Interview

It would not be surprising if individuals are confused, and do not understand, that HIPAA [the Health Insurance Portability and Accountability Act] may not protect the privacy and security of their health information collected by equipment or an app if that collection of information is not offered by the individual's provider or on its behalf... [S]ome mHealth developers themselves may not be aware of the regulatory requirements that attach to their work and have requested additional guidance.[98]

In the domain of corporate wellness programs, privacy and identity expert Anna Slomovic cited the issue of sensor fusion and the general promiscuity of data, observing that the US's sectoral approach to privacy may be insufficient. She argued in favor of a broader approach to privacy regulations:

*Wellness programs give us a glimpse of a future where data collected from multiple devices and multiple sources will be combined and analyzed, resulting in a level of surveillance that is much greater than the sum of its parts... **While US privacy protections are sectoral, data flows in the real world are not.***

As more objects get connected to the Internet, it will be more and more difficult to confine their data within a single regulatory silo. *Current tools will be utterly inadequate in the new connected world.*[99]

In IoT discussions, several stakeholders and researchers seek the creation of an omnibus regime with baseline privacy rules.[100] When asked about the efficacy and validity of the reasonable expectation of privacy standard, a fundamental element of US privacy governance, Michelle De Mooy responded:

I think it's poor. It's definitely not ideal, but at the same time, when you don't have a baseline privacy law it makes it incredibly difficult to not have that standard. The expectations are set in some ways by the laws that are sector-specific, and so they perhaps have shaped public perception of what they consider sensitive, but I don't think there's any way around that. I wish that we had some sort of baseline legislation or law that would give people a different set of common, everyday expectations for their privacy.[101]

98 US Dep't of Health and Human Services, 2016, p. 9-10

99 Slomovic, 2015, emph. added

100 Comparisons of the merits of sectoral versus omnibus privacy rules have been debated by numerous other scholars. See, e.g., Schwartz, 2013

101 Interview

Heather Patterson of Intel echoed Anna Slomovic's concerns for the inadequacy of a sectoral regime in the face of context blurring:

When it all becomes one broader context and there are vast treasure troves of information about us that can be assembled and reassembled, then we need to be mindful about whether omnibus privacy protections are now warranted and whether we want to up-level privacy regulations so that it follows people [everywhere] rather than the spheres that the people happen to be functioning within when that data becomes collected from them.[102]

Law professor Paul Schwartz observed that the United States is an anomaly in its absence of baseline privacy rules: "The United States' unique path as a matter of form (no omnibus law) and substance (a limited set of [fair information principles]) has made it an outlier in relation to the global community."[103]

In its 2015 IoT report, the FTC reiterated its call for "broad-based (as opposed to IoT-specific) privacy legislation."[104] The Department of Commerce concurred, stating that it "continues to address privacy concerns in a range of contexts, from support for baseline privacy legislation that would include IoT services, to work to promote the availability of strong encryption (including in IoT devices)."[105]

Many industry stakeholders, however, oppose changing the policy landscape. Comments submitted for the aforementioned Dep't of Commerce IoT paper by commercial companies and trade bodies are illustrative:

Although IoT enables the collection of data through inter-connected devices that previously may not have been feasible, the underlying data collection is not a new policy issue that would warrant a new framework for addressing privacy... The FTC's existing privacy framework, which incorporates many of the Fair Information Practice Principles (FIPPs), is well suited to address the IoT environment.[106]

102 Interview, emph. added
103 Schwartz, 2013, p. 1636
104 Federal Trade Commission, 2015, p. viii
105 US Dep't of Commerce, 2017, p. 42
106 Epps, 2016

*Self-regulation is nimble, and can be more easily updated to address changes in the marketplace and technology... In fact, self-regulatory codes may be the **best** way to effectuate consumer adoption of the IoT. As a backstop with respect to consumer privacy, the FTC can utilize its Section 5 authority to protect against any privacy-related practices that are unfair or deceptive.*[107]

Such views comport with the commercial logic that increases in regulation cause increases in costs and limitations on product developments.

Despite the US's top privacy regulator and the agency tasked with multistakeholder engagement on IoT publicly supporting baseline privacy regulation, prospects seem dim. One prior legislative attempt, the 2015 Consumer Privacy Bill of Rights Act, was stillborn,[108] failing to galvanize consumer protection groups[109] or gain bipartisan support. Omnibus federal privacy legislation has yet to reach an advanced stage in Congress, and the deregulatory character of the current Administration in combination with a Republican-led Congress points to a reduction in the already low chances of such legislation occurring.

The European General Data Protection Regulation (GDPR) coming into force in May 2018 is a substantial upgrade to the EU's existing omnibus data protection rules. Its impact will affect American companies since the Regulation applies to entities processing Europeans' data, irrespective of a company's geographic location. Compared to the existing US privacy regime, the GDPR requires far more internal assessment and compliance of data practices, and it threatens sanctions for non-compliance. It will test the argument that increased regulatory burdens can stifle innovation.

107 Kearney and Reynolds, 2016
108 See Singer, 2016, for one account of the causes
109 See, e.g., an open letter from numerous groups to the White House, available at http://www.consumerwatchdog.org/resources/ltrobamagroups030315.pdf

THE NEED FOR ENHANCED SOCIAL DISCOURSE

A major theme to emerge from interviews and our two workshops was a perceived need for more social conversations about the collection and use of personal data, in tandem with more transparency on the part of data collectors. This theme reflects an awareness that public understanding of the coming waves of IoT devices, essential to informing governance and the evolution of norms, is currently inadequate. Lee Tien of the EFF linked the need for transparency on the part of data collectors to society's ability to hold this conversation:

If we have the policy knowledge about what is happening to our data – how well it's being stored, how well is it being de-identified, what sense and meaning can be derived from it at this present level of de-identification – we'd be able to get more of the information that we need to make a social, public decision about what the ideal is... That's one area that I think the law is very important; why [advocates] often focus on transparency and disclosure. Consumers should be able to know where their data goes. They should be able to know, if you claim you are de-identifying data that you collect, what is the protocol that you use for de-identifying? What data is that FitBit actually collecting? Where does it actually go? Who are your partners?[110]

Dawn Nafus, an anthropologist at Intel, argued for society's need for a greater capacity to critically understand data: "If more people had a clear idea of what data actually meant and what it doesn't mean… then you can start to deepen the critique and the resistance against the bad uses, as well as elaborate the good ones."[111]

110 Interview
111 Interview

Tien expressed optimism about the evolving state of IoT discourse, remarking that interest in privacy and security has gained mileage in public discussion over the last decade or two.[112] However, Heather Patterson had a less sanguine view on the state of discourse on IoT data and privacy:

We need more social discourse about what is happening right now, and I'm just not seeing it there and [instead I'm] seeing lots of conversations about shareholder value. Frankly, I just think the world is bigger than that.[113]

These comments are not restricted to the IoT – they fit comfortably in conversations about big data, or indeed about the generally datafied world we now live in. Fortunately, privacy remains a popular topic among the general public and specialists alike. As Nafus and Tien note above, deepening the capability for critique is a part of reaching a better ideal of the 'information society.'

The converse of this need to have a social conversation about data use is a general lack of understanding about connected devices and technology. Interview and workshop participants noted that ignorance pervades the offices of policymakers, politicians, their lawyers, and the public. Rafi Rich, a smart city consultant and former director of buildings and licensing for Israel, remarked on the privacy issues surrounding the deployment of smart city technology:

Most mayors don't have a clue whatsoever about the dangers of data going around, so they basically don't see it as a problem, because although there are regulations around for privacy, ***most cities that I have been working with don't really understand that when they give services to private companies there is a risk.*** *The local authority... I won't say they don't care, but they just don't have enough knowledge to care. I think that local authorities in most places around the world are still in the 19th century. Technology moves much too fast for them.*[114]

112 Interview
113 Interview

114 Rich, Interview

Speaking of his experiences while principal advisor on technology, data and privacy in the California State Attorney General's office, Justin Erlich lamented the lack of in-house technical expertise and availability of external help:

I saw our lawyers wanting to do the right thing with consumer protection and the internet, but the problem is: 1) there is no clear law on emerging tech and lawyers can be risk averse, 2) there's a need to get creative about legal theories, but we don't completely understand the underlying technology and how it works so it's hard to develop them... It feels like bringing a knife to a gunfight.[115]

Lee Tien expressed his view of government policymakers even more bluntly:

My view is that they don't know technology well enough to know when they don't know... I have engineers who brief me on technology I don't know; I probably have more technical support than legislators and policymakers. That's wrong. ***Why do I have more access to more cryptographers, machine learning experts, and technologists than people who actually make decisions?***[116]

Greater social comprehension of data use is a strategy that nearly all IoT governance participants seek. Transparency in the form of privacy policies is encouraged by industry and the privacy community alike, and both sides see benefit in people being much more informed about how data is collected, analyzed, and shared. We discuss different ways to encourage greater awareness in the *Emerging Frameworks* and *Recommendations* sections.

115 Workshop comment

116 Interview, emph. added

CONCLUSION

The question as to whether the IoT is different enough from prior internet technologies to warrant its own special regulations is in no way settled. This is partly because there are stakeholders on both sides of the discussion with much to lose – an increase in regulation will likely result in higher compliance costs, but an absence of appropriate regulation could further injure the state of privacy. As the Dep't of Commerce paper cited above reports, policy stakeholders have yet to come to a consensus that there is a privacy problem.

IoT governance will proceed in much the way that policy proceeds generally: incrementally, sometimes haphazardly, with greater and lesser degrees of order and coordination. The difficulty of creating adequate privacy rules is exacerbated by the multiplicity of agencies who claim (and refuse) jurisdiction. US state legislation will continue to press onward based on its own institutional dynamics, sometimes in tension with federal prerogatives. IoT policy is inseparable from politics, history, institutional arrangements, and power, and so the choice to regulate different aspects of the IoT ecosystem – data gathering and use, security, liability – will inevitably be influenced by existing rules, the power of stakeholders, dominant economic philosophies, discourse, and perhaps the infrequent, unexpected 'shock' to the system, such as a massive denial-of-service attack.

The GDPR represents the most ambitious information policy project on the globe. It envisions affecting Europeans and non-Europeans alike, and is an attempt to export European privacy and data protection values to wide range of countries, some of whose companies (read: America's) will inevitably be resentful of it. It will test whether sectoral or omnibus regulations are better suited to address the privacy challenges connected devices, and as such will provide an analytically rich counterpoint to US IoT policy evolution.

The mechanics of public policy are a rarified subject, but public discourse remains an important input to policymaking. As our research respondents have highlighted, there is value in arming the populace and policy actors with more information to consider IoT privacy issues. For the general public, enhancing the transparency of data collection and use practices is one important method, as are the adversarial and educational approaches of advocacy groups and journalists. In policymaking itself, there is a need for more technologist interaction with policy actors, which will only occur if the labor economics are well-aligned.

PRIVACY MANAGEMENT

The Internet of Things heralds a qualitative shift in the way privacy is managed, both by people and by the organizations that create, sell and operate connected devices. It amplifies prior privacy challenges like the opacity of data flows and actors, and it creates new issues, such as creating a false sense of trust due the familiarity and unobtrusiveness of devices.

This section explores the IoT's effects on individuals' ability to manage their own privacy – issues of awareness, understanding, and control. In particular, we discuss the IoT's impact on transparency, notice and disclosure, choice and trust. We explore the significant privacy impact of forgetting that data-collecting objects are present in our environment, and delve into people's changing relationships with their devices, in particular the impact of the IoT on children.

THE IoT'S IMPACT ON TRANSPARENCY

Any attempt to manage privacy presupposes that people are aware of the collection, use and disclosure of their personal data by others. Actors in privacy-sensitive arenas are generally expected to practice transparency in order to facilitate this awareness. The 1973 US government report "Records, Computers and the Rights of Citizens" envisioned five Fair Information Practices, two of which state:

- **There must be no personal data record-keeping systems whose very existence is secret.**

- **There must be a way for an individual to find out what information about him is in a record and how it is used.**[117]

This report directly informed the US Privacy Act of 1974, America's only omnibus privacy legislation;[118] the modern Fair Information Practice Principles (FIPPs); and the OECD's 1980 Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, which in turn directly informed the EU's 1995 Data Protection Directive and the recently enacted General Data Protection Regulation (GDPR).

The US FIPPs call for Transparency: "[An organization] should be transparent and provide notice to the individual regarding its collection, use, dissemination, and maintenance of personally identifiable information (PII)."[119] In the GDPR, transparency is a key part of ensuring fairness, intrinsic to consent, intended to go beyond privacy notices, and is intended to be "user-centric."[120]

117 US Dep´t of Health, Education and Welfare, 1973

118 Though it only applies to Federal entities.
119 US Dep't of Homeland Security, 2008
120 Center for Information Policy Leadership, 2017

The Transparency principle and its antecedents stem from the democratic principles of autonomy and self-determination: a person must be aware of what data is collected about them and how it's used to fully construct herself and act freely. Transparency about data collection and use enables people to make informed choices about the direction of their informational lives. It also underpins other privacy practices, such as the ability to act upon, correct or delete data, and the ability to exercise a measure of control over the information central to one's life. Connected devices are intended to be unobtrusive, to blend in with the environment; and in many cases they take the form of familiar objects that have been upgraded with networking, sensors, and other new functions. Wearables look like watches, jewelry and fitness gear. Amazon Echoes look like speakers. Hello Barbie looks like mute and deaf Barbie. Voice-activated televisions look like… televisions. In late 2017, Amazon released the Echo Spot, which looks like a clock radio with a large screen and, based on its marketing, is intended to be put on the nightstand in a similar way.[121] The difference is that it has a small camera at its top. And, while clock radios do not gather data, this one explicitly does, combining it with artificial intelligence and sending it back to Amazon to assist in the compilation of detailed profiles about people's preferences. This camouflaged monitoring of our activities puts the IoT in tension with Transparency. A 2016 report by Canada's Office of the Privacy Commissioner (OPC) summarized the issue:

Amazon Echo Spot

As these technologies become ubiquitous, we may have little or no warning or awareness that they are even in place; ***they simply become part of the backdrop of our daily lives.*** *How, then, can citizens who may or may not want to use this technology ensure that someone is held accountable for its use? How will they be able to challenge how the information is used, and how will they be able to give any kind of meaningful consent?*[122]

121 See Introducing Echo Spot: https://www.amazon.co.uk/Introducing-Amazon-Echo-Spot-Black/dp/B01J2BK6CO

122 Office of the Privacy Commissioner of Canada, 2016, p. 23 emph. added

The unobtrusiveness of connected devices means that data collection activities are invisible. To be sure, invisible data collection is the hallmark of the internet era – from the logs of the first computer networks, to tracking on the web of the 1990s, to modern, mature online profiling, to the proliferating IoT devices of today, data collection has happened in the background, silently. Discussing the ethics of online behavioral tracking for advertising ends, privacy expert Omer Tene noted in 2011 that:

> *Users are seldom aware of the data collection processes, prospective data uses, and identity of the myriad actors involved, including not only advertisers and publishers, but also ad networks, ad exchanges, analytics services, affiliate marketers, market researchers, and search engine optimizers... As a result, behavioural targeting clearly challenges the [OECD's and EU Data Protection Directive's] principles of transparency and individual choice.*[123]

The IoT amplifies these challenges in its expansion of data collection from the online to the offline world. This is all the more true when one considers data collection by other people's devices and IoT devices in public spaces.

THE IoT'S IMPACT ON NOTICE AND DISCLOSURE

Online data gathering is subject to norms and laws of disclosure. Privacy policies for websites and applications are institutionalized and often mandated, though the efficacy of providing notice about data practices has been widely debated.[124] However, in these early years of IoT development and device sales, research reveals a diminishing adherence to these norms. The issue is two-fold: 1) the reduction or elimination of screens from IoT devices makes the display of privacy notices problematic, and 2) some device manufacturers are being lax in their duty (legal or norm-derived) to disclose data collection and use practices.

123 Tene, 2011, p. 5

124 See, e.g., Calo, 2012; Sloan and Warner, 2013

The loss of screen space and resulting impediment to displaying privacy policies has been discussed by a number of scholars and stakeholders. The FTC stated in its 2015 IoT report:

> *Many IoT devices – such as home appliances or medical devices – have no screen or other interface to communicate with the consumer, thereby making notice on the device itself difficult, if not impossible. For those devices that do have screens, the screens may be smaller than even the screens on mobile devices, where providing notice is already a challenge. Finally, even if a device has screens, IoT sensors may collect data at times when the consumer may not be able to read a notice (for example, while driving).*[125]

However, Michelle De Mooy of the Center for Democracy & Technology was optimistic about notifications improving in the IoT space: "When you view transparency as integral to the functioning of a device or product or service, the opportunities for engaging with users become more broad and innovative."[126] Meg Leta Jones of Georgetown also sees opportunity for a re-think of notification in the IoT: "The IoT just strips away the screen and strips away the personal computing aspect of privacy and offers us this new platform to reconsider what privacy is and how we want to manage it."[127]

As to IoT manufacturers being lax in disclosures, a pattern of poorly informing users is emerging. Prof. Scott Peppet surveyed twenty popular IoT devices in 2014, including the Nest Thermostat, the FitBit, health products, and home monitoring systems, in an attempt to gauge the depth and degree of their privacy disclosures.[128] He found them to be shockingly inadequate:

> ***None*** *of the twenty devices included privacy- or data-related information in the box. None even referred in the packaging materials or user guides to the existence of a privacy policy on the manufacturer's website... Some policies seem to apply to both website use and sensor-device use. Other policies limit their application to website use, not sensor-device use, but provide no means to locate a device-related privacy policy. This leaves unanswered whether* ***any*** *privacy-related policy applies to the data generated by these devices.*[129]

125 Federal Trade Commission, 2015, p. 22
126 De Mooy, Interview, emph. added
127 Interview
128 Peppet, 2014
129 Peppet, 2014, p. 140-142, orig. emph.

Europe's Article 29 Data Protection Working Party (Art29WP) has taken a strong position against vague or general notices: "Information and consent policies must focus on information which is understandable by the user and should not be confined to a general privacy policy on the controllers' website."[130] But Peppet found the language in the devices' policies to be exceedingly ambiguous:

These policies are often confusing about whether sensor or biometric data count as 'personal information' and thus unclear about how such data can be shared with or sold to third parties... privacy policies for consumer sensor devices often do not mention ownership of sensor data. Of the twenty products covered... only four discussed data ownership explicitly... only three provided clear information on exactly what sensors the product included or what sensor data the product collected... many of these Internet of Things privacy policies provided no information on what sensor data their device generated.[131]

In the same vein, a 2013 study of 23 paid and 20 free mobile health and fitness apps found that:

— 26% of the free and 40% of the paid apps had no privacy policy

— 39% of the free and 30% of the paid apps sent data to someone not disclosed in the app or the privacy policy[132]

And, in 2016 a study of over 300 devices by 25 of the world's data protection authorities found:

— 59% of devices failed to adequately explain to customers how their personal information was collected, used and disclosed

— 68% failed to properly explain how information was stored

— 72% failed to explain how customers could delete their information off the device

— 38% failed to include easily identifiable contact details if customers had privacy concerns[133]

130 Article 29 Data Protection Working Party, 2014, p. 22
131 Peppet, 2014, p. 142-144, orig. emph.
132 Ackerman, 2013, p. 19-20
133 UK Information Commissioner's Office, 2016

Some of our respondents believed that companies were intentionally opaque about their data gathering and use practices. Michelle De Mooy remarked:

Everything is stuffed into a privacy policy that of course no one reads, and they know that. I think it illustrates how important trust truly is and the fact that companies do understand that, because they hide so much of what they're doing in either doublespeak or in these lengthy privacy policies or terms of service that they know that their users aren't likely to look at. So, I think it's a deliberate opacity on the part of the industry.

Lee Tien of the EFF shared this view:

It's all designed to create more of a Potemkin Village of disclosure, but nothing that actually allows you to say, "Wow, that's all going to Company X? Maybe I should be worried that Company X is getting all of this information from the 79 different companies that I do business with."[134]

However, it is difficult to know whether this opacity is deliberate or whether it is simply impossible to write intelligible privacy policies that are legally cognizable and adhere to regulation, an issue we address in the *Governance* section. Strong research exists that the general level of privacy policy comprehension is low,[135] and the device surveys we cite above illustrate an alarming trend of diminishing notification quality in the IoT. But it is still early days, and there are serious research attempts to address the above shortcomings, which we detail in the *Emerging Frameworks and Strategies* section.

134 Interview

135 McDonald and Cranor, 2008

FORGETTING DATA-COLLECTING DEVICES ARE NEARBY

The combination of diminished screens and lax disclosure of privacy information makes it hard for purchasers of IoT products to understand what these devices see, hear and know, and how their manufacturers and other parties will use the collected data. This issue is at the heart of the Canadian OPC's observation above: "We may have little or no warning or awareness that they are even in place; they simply become part of the backdrop of our daily lives." Ergo, people will forget about the presence of data-collecting devices, and there are privacy management risks from forgetting devices are around. Elaborating on the relationship between behavior and IoT monitoring technology, Florian Schaub of the University of Michigan remarked:

You think you're being observed, so you behave differently, but there are also habituation aspects. You might behave differently at first, and then you become attuned to the technology being there, so you become maybe a bit more laissez-faire in terms of how you behave. ***You forget that these devices are active, but that doesn't mean that you're not leaving digital traces of mundane behavior anymore****, and this data can reveal information about your preferences, what you like, what you don't like, or your health, your family situation, your financial situation – all kinds of different things that people prefer to keep private.*[136]

Meg Leta Jones saw such habituation as partly a result of anthropomorphizing our devices:

We can be duped a little bit by our objects. *We can get really comfortable with them really fast; anthropomorphize them in ways that cause us to divulge information beyond what we would intend.*[137]

136 Interview, emph. added
137 Interview, emph. added

Law scholar Margot Kaminski remarks upon the same problem in her analysis of the privacy implications of home robots:

There is evidence that people treat anthropomorphic robots with increased compassion and trust. A robot that lulls people into revealing more than they intend to may be viewed as deceptive technology; or it may be treated similarly to false human friends.[138]

In her anthropological research on IoT devices users, Heather Patterson of Intel hears of discomfort occurring at the boundary of forgetting and awareness:

We forget that they're on until they wake up and surprise us, or let us know that they've been tracking us all along by telling us how long it will take us to get to work before we've even asked. We realize, "Oh, my phone completely knows... Google knows my pattern and I wasn't aware of that." They're intrusive.[139]

This discomfort serves as a reminder that, most of the time, people are simply not fully cognizant of the level of surveillance going on around them. Addressing this, L. Jean Camp and her colleagues argue for the "engineering design principle of 'least surprise': The IoT should behave in a manner that is both expected by and clearly communicated to every stakeholder with which it interacts."[140]

Patterson notes how the problems of forgetting about the devices are likely to be exacerbated by their further integration in human environments:

It's not just about IoT anymore: it's IoT plus artificial intelligence plus machine learning. We're now facing the prospect of using devices that are always on, always watching, always listening and always talking to one another. Particularly in closed spaces like the home or the workplace, I think we're looking at a future where there's really a strong possibility of those systems being embedded in our built environments.[141]

138 Kaminski, 2015, pp. 666 and 671
139 Interview
140 Camp et al., 2016
141 Interview

Relatedly, workshop participants raised the idea that connected device owners may have poor mental models of how they gather data. Jen King, Director of Privacy at Stanford's Center for Internet and Society, noted that her research on RFID technology revealed that users universally expected some kind of feedback when data was being collected.[142] More generally, King saw two kinds of products emerging: "new products that we will have never encountered, and everyday products that have IoT functionality slapped on. We already have well-established mental models about the objects around us, so expecting users to make that kind of leap or understand the tradeoffs is unproductive because they have these set mental models."[143] Ergo, everyday products that acquire IoT functionality put an increased strain on people's already nebulous understanding of data harvesting.

Reduced knowledge about personal, intimate data collection; decreasing attention to the presence of monitoring devices; and forgetting devices are around while being lulled into revealing ever more personal information all contribute to a **diminishment of private spaces**, a topic we discuss in greater detail in the *Risks of Harm* section.

Reduction of awareness and understanding of privacy risks affects another aspect of privacy management: Choice, the ability to *knowingly* accept or decline a data-gathering product or service, which is a mainstay of privacy regimes. Florian Schaub noted that our freedom to choose is increasingly limited to just the initial purchase of a device, with little control over what happens afterward:

> *In the home environment you don't really have that much control over your privacy with IoT devices. Your biggest control element is deciding which devices you place in your home and vetting them for good privacy practices.*[144]

Schaub also notes, "It's often difficult to find this information for consumer devices and take it into account in any kind of purchasing decisions," which is supported by the research we discuss above regarding the inadequacies of IoT privacy notices.

142 Workshop comment
143 Workshop comment
144 Interview

Importantly, the Canadian OPC argues that the ability to choose erodes as IoT functions become more prevalent and 'non-smart' devices are less available:

As smart devices and appliances become more and more normalized, there is an increasing 'erosion of choice' for individuals who would have preferred their 'non-smart' versions.[145]

These are serious challenges to reliance upon Choice as a fair information principle. **Market shifts towards 'smart' features that are intentionally unobtrusive lead to less understanding of data collection, and less ability to decline those features.** Choice is further eroded as sensors, cameras and microphones become standard features. If all televisions have cameras, how can one choose not to have a camera staring into the faces of one's family when watching a movie? The problem is further exacerbated by the presence of devices we do not own or control – the 'Internet of Other People's Things,' a term used by Meg Leta Jones[146] to describe the growing number of devices that not only capture sensitive data about their owners, but also the people around them.

145 Office of the Privacy Commissioner of Canada, 2016, p. 21, emph. added

146 See Jones, 2014

PEOPLE'S CHANGING RELATIONSHIPS WITH THEIR DEVICES

Slowly but surely, smart devices are becoming entire smart environments, homes and public spaces. Things that were inert are becoming aware, fundamentally changing the way people interact with and relate to their physical environment. A TV is no longer just a broadcast device – it is a transmitter/receiver/observer. A toy is not simply a canvas for a child's imagination – it is a puppet whose strings are pulled by far-off processes, seeking self-improvement via the data it gathers. A speaker is a virtual assistant with, as technology law scholar Ryan Calo puts it, a "tiny salesperson"[147] inside, fusing your real-time lived experience with past purchases, preferences, and predictions.

147 Calo, 2013

These devices are not neutral; they are constructed with a commercial logic encouraging people to share. **The IoT embraces and extends the logic of social media – intentional disclosure, social participation, and continued investment in interaction.** There is no doubt that disclosing data in new ways is fun and valuable to the purchasers of devices, but networked, sensing consumer products have special characteristics that must be investigated and critiqued.

Children's use of the IoT is an instructive case. One of the most visible and known examples of internet-connected toys is Hello Barbie.[148] Introduced in 2015, this upgraded version of the nearly 60-year-old Barbie doll connects to the cloud and is able to parse children's language and talk back. Before this modern automaton was even in the stores it stoked controversy, drawing the ire of a group called Campaign for a Commercial-Free Childhood (CCFC), who campaigned against its sale under the banner "Hell No Barbie."[149] CCFC opposed children's conversations with the doll being shared with corporations who have a commercial interest in them, believed that parents should not have access to the recordings, and argued that Hello Barbie undermines creative play. Hello Barbie's technical advancements – speech recognition, cloud-based processing, and easy Wifi access – have given rise to new relationship arrangements between children, objects, their parents, and the myriad stakeholders who bring these new toys into existence.

The new relationship dynamics introduced by Hello Barbie's features raise serious trust issues, especially as its young users are probably unaware that Hello Barbie is recording everything they say. Jones and Meurer observed:

148 See http://hellobarbiefaq.mattel.com/

149 Campaign for a Commercial-Free Childhood, n.d.

Hello Barbie frequently requests the trust of both parents and their children, yet it is simultaneously built to undermine those relationships of trust. It is superficially designed to act as a child's best friend, but the doll records and shares every conversation with a child's parents. Parents, on the other hand, are given the perception of complete control over the stored conversations, yet their child's data has the potential to be shared with numerous third parties... If Hello Barbie were a real person, she would be more likely to be known as the Gossip Queen than a trusted friend.[150]

Internet-connected toys are a useful case by which to illustrate how IoT devices affect relations between people. Particularly, **the burden of managing children's privacy has historically fallen upon parents. The IoT increases the weight of that burden.**

CONCLUSION

The Internet of Things is catapulting us into a new era of surveillance. IoT devices can blend into the background, take on the form of a child's best friend, or encourage people to willingly share their data on social networks. Over time, connected devices will become so well-integrated into our lives that we forget they are there. Meanwhile, the nature of IoT devices makes it difficult to practice transparency and implement easily accessible privacy policies, especially in the case of screenless IoT products, rendering it very difficult to make informed choices regarding these devices. Early evidence suggests that IoT manufacturers are making this problem worse by being lax about the breadth and availability of privacy notices.

Still, the new relationships created by technology evolution can bring new benefits, new forms of entertainment, and new, valuable ways for families to interact. Rather than reject these new devices and relationships outright, we wish to highlight their privacy deficiencies and poorly understood data-gathering qualities in service of better policies, product designs and norms.

150 Jones and Meuer, 2016, p. 4

INCENTIVES FRAMEWORK AND MECHANISMS

Understanding organizational incentives is key to identifying governance mechanisms that can address IoT privacy issues such as choice and personal data management, and informing policies that encourage appropriate privacy protections.

These incentives are shaped by the economic dynamics of IoT privacy. Some of the factors that contribute to the current lack of market- and regulatory-based incentives for industry stakeholders to embrace IoT privacy management include:

— Insufficient information about the costs and benefits of stronger or weaker privacy protection to different IoT stakeholders;

— Confusion about data ownership and responsibility across IoT data flows and supply chains, including challenges with tracking data provenance from collection to use in order to monitor and enforce privacy in the IoT ecosystem;

— Lack of technical understanding and awareness, impeding overseeing entities from asking the right questions and seeking effective solutions;

— Uncertainty about IoT liability exposure, and awareness of supply chain vulnerabilities;

— Inadequacy of the collective-action mechanisms operating as forcing functions for privacy management – for example, data breaches have become normalized and consumer breach fatigue has set in, so industry motivation (e.g., fear of consumer lawsuits, reputation damage, regulatory action) has diminished;

— Lack of incentives for manufacturers to invest in adequate security measures, given the low cost and/or low margin nature of IoT devices;

— Diminished accountability, as the majority of relationships in the IoT ecosystem are business-to-business and not business-to-consumer.

One framework that can serve as a model to understand and improve incentives for IoT privacy is the ensemble of methods which have been established to combat environmental pollution, as the two issues share some similar features.[151] [152] In this section, we analogize the forces that induce a reduction in the release of harmful pollutants to the mechanisms that incentivize the abatement of collection, use or disclosure of privacy-sensitive

151 See, e.g., Hirsch, 2006; Froomkin, 2015; Foulon et al., 2002

152 Note that the range of potential incentive mechanisms we offer generalizes to any resource management system. Here we offer up privacy-inducing information as the resource to be managed and encourage people to think expansively about the application of these incentives to IoT privacy.

information ('personal data emissions') that infringe on privacy ('privacy pollution').[153]

In general, mechanisms that aim to manage the consumption and production of data that compromises privacy rights and interests fall into two broad categories. The first encompasses regulatory approaches that set standards for collection, use and/or disclosure of privacy-sensitive data, while the second comprises pure market-based approaches that rely on economic forces to correct for producer and consumer behavior. In general, regulatory approaches provide organizations with more certainty at the expense of flexibility, while the market-based approaches do the reverse. There are also several variations of these approaches, including hybrid approaches that combine elements of the two.

A regulatory standards approach modeled after environmental regulations could set personal data emissions limits in order to reduce externalized costs to the IoT ecosystem and broader internet, but may impose potentially large costs on polluters via fines and litigation, as well as enforcement costs on victims and oversight entities.

In contrast to regulatory standards models, market-based approaches can be price- or quantity-based and include mechanisms such as: permit/credit/cap-and-trade systems; taxes, fees, and charges; subsidies; tax-subsidy combinations; self-regulation, insurance, and revenue/remuneration. Privacy Reduction Credits would set limits on rates of personal data emissions, whereby companies could earn credits by reducing privacy pollution emissions. A cap-and-trade variant would set a maximum allowable cap on total personal data emissions. Both would allow companies the choice between reducing their privacy pollution emissions or purchasing allowances from low privacy polluters. The tax-based approach would place a per-unit monetary fee on privacy pollution. The disadvantage is that this approach would not guarantee a specific amount of reduction in privacy pollution, as it limits itself to penalizing violators after the fact. A subsidies mechanism would reward privacy polluters for reducing personal data emissions and could take the form of grants, low-interest loans, favorable tax treatment, or procurement mandates/preferences. Through this scheme, the accrual of benefits would encourage companies to enter the market.

153 US Environmental Protection Agency, n.d.

Hybrid approaches create regulation using a mix of standards and pricing combinations, liability rules, and information disclosure. A standards-pricing mechanism would create a personal data emission standard that imposes a tax for emissions that exceed the standard. Liability law (negligence, strict liability, breach of warranty) is one of the logical candidates for application to the IoT, but it is unclear to what extent it could support harmful IoT personal data emissions claims, if at all.[154] Disclosure regimes include both mandatory reporting, such as state data breach laws, and voluntary reporting, as in labeling schemes like the proposed EC Trusted IoT Label[155] or the IoT Security Foundation's Best Practice User Mark.[156]

Voluntary market-based approaches are currently the prevailing force at play with regard to IoT privacy in the US, although government regulation is gaining ground in specific sectors such as connected vehicles, unmanned aerial systems, and medical devices.[157] Market-driven approaches are predicated on the promise that the market will reward companies who reduce or eliminate personal data emissions by providing them with a competitive and/or economic advantage. The purported benefits of voluntary approaches for industry would be largely reputational (improved public image, resulting in customer loyalty) and/or preemptive, in that self-regulation can inform and ease transition to formal law by amortizing the costs associated with becoming compliant when law does go into effect.

154 Modern cyber-risk contexts characterized by software defects and insecure devices expose loopholes in the traditional product liability regime – strict liability, negligence, breach of warranty. This is primarily due to the economic loss doctrine which bars recovery for productivity loss, business disruption, and other common damages caused by software defects. As well, the application of design defects principles to software is difficult given the complexity of the devices and recent tort reform trends that have limited liability. Further, the intervening cause of damage from insecure software is typically a criminal or tortious act by a third party, so principles of causation might limit liability for manufacturers. See, e.g., Dean, 2018.

155 European Commission, 2016b

156 IoT Security Foundation, n.d.

157 See, e.g., AV START Act, S.1885; Federal Aviation Administration, Unmanned Aircraft Systems, JO 7200.23A

COUNTERING MARKET FAILURE OF PRIVACY

It remains to be seen whether privacy in the IoT will become a market failure that results in costs to society and necessitates stronger government intervention.[158] The question thus arises: left to its own devices, would the IoT market fail to find an appropriate degree of privacy protection for consumers? In cases of market failure due to information asymmetry, transparency and disclosure of privacy practices may be of use in informing more optimal decisions for industry and consumers when attempting to align incentives to counter a possible market failure of IoT privacy. Market-based or hybrid approaches that force companies to internalize the costs of personal data emissions would be possible strategies that could help correct for negative externalities.

It is generally agreed that there is a lack of incentive for transparency of data flows. Lee Tien of the EFF noted:

> *We know that the companies and the law enforcements and governments [don't] have a strong institutional incentive to tell everybody what they're doing. Neither does Google... That's a gap that hopefully law can fill.*[159]

One example of a disclosure and transparency mechanism would be a rating system to correct information asymmetries, foster the IoT privacy controls market, and socialize the value of privacy attributes of the IoT. Consumer Reports´ Digital Standard is one possible vision of this.[160]

If privacy is to evolve into a market differentiator, there must be clear agreement as to when personal data collection, use and disclosure are required in order to deliver services and ensure core functionality without unduly impeding the utility of the IoT.

Another possible market-correcting mechanism along the transparency continuum would involve presenting the trade/sale/exchange of personal data to

158 See, e.g., President's Council of Advisors on Science and Technology, 2014; Camp and Johnson, 2012; Schneier, 2017; Vila, et al., 2003

159 Interview

160 See The Digital Standard https://www.thedigitalstandard.org/

consumers as an explicit and transparent bargain. Currently the bilateral value exchange where users create profiles and provide data in exchange for free or low-cost products and services is assumed without any real dialogue or conscious choice by users. Tel Aviv's DigiTel Resident Card serves as an example of explicit exchange, wherein users give consent for the government to leverage existing IoT infrastructure.[161] Companies then tap into the existing behavioral sensors (e.g., location data on mobile devices), ultimately resulting in cost savings that are passed on to users.

Regulation is not inherently at odds with market competitiveness and innovation: markets do not and cannot function without rules. On the one hand, technology-neutral regulation has the advantage of being more generalizable and future-proof. However, there is the risk it will fail to serve as a forcing function due to the discretionary latitude in enforcement. On the other hand, technology-specific regulation, while providing more certainty to the market ex ante, lacks the flexibility that is needed when dealing with developing technology that does not yet have accurate cost-benefit assessments. The IoT market represents a confluence of burgeoning social and political factors.

The GDPR may serve as a prototype regulatory incentive mechanism due to its strong monetary penalties for failure to properly manage personal data.[162] As well, it may produce positive externalities for the regulated entities, such as economic benefits to business operations by way of enhanced efficiency. For example, it could provide homogenous data governance between a connected car (device) and a smartphone (service).

A variant of the regulatory approach would involve stronger mechanisms to help victims of data breaches recover monetary damages. While the GDPR affords private rights of action, most current US privacy laws do not, relying instead on government enforcement actions or preempting litigation opportunities with mandatory arbitration. Neither approach gives people any incentive to hold firms to account. As well, regulation or common law could focus on placing more responsibility on intermediary platforms, either by mandating some form of privacy protection or by enforcing transparency in the ways data is collected and used. It could also take the form of clarifying data breach notification requirements to cover breaches of customer IoT data. These solutions, while also imperfect, have the benefit of

161 See Resident's Card: https://www.tel-aviv.gov.il/en/Live/ResidentsCard/Pages/default.aspx

162 The GDPR can sanction up to 4% of global revenue or up to €20 million.

sidestepping the issue of quantifying harm, which both taxes and private rights of action would probably require.[163]

Another way to address negative externalities is by directly taxing companies that produce privacy pollution (do not comply with privacy standards/produce personal data emissions). Rather than banning privacy pollution outright, each privacy polluter could choose their preferred means of compliance, such as in cases of firms choosing whether to comply with the 1990 Clean Air Act or pay a carbon tax. Companies could then raise prices, implement privacy protections, reduce the production of privacy-sensitive data, or choose to pay for their violations. The challenge to this approach lies in the difficulty of quantifying the size and cost of privacy harms. In both cases, there is a behavioral economics aspect of IoT privacy risk: considering consumers' tendency to underestimate the severity, likelihood and timeframe of IoT privacy risk, regulations would need to be scrupulously designed in order to help consumers better gauge these risks.[164] For example, users' sense of privacy risk from willingly sharing seemingly benign location or sentiment data may present as short-term, small severity and low likelihood risks of harm. Users are unlikely to realize that collection over time and aggregation with other such data may result in privacy violations that are high likelihood and large severity. Regulation would do well to address these risk factors in prescribing industry practices.

The transfer of privacy risk vis-à-vis insurance is another mechanism to incentivize companies to address IoT privacy risk, although it can carry moral hazard concerns when companies wholly disregard their privacy pollution under the mistaken belief that insurance will unequivocally address privacy liability. US commercial liability insurance for privacy harm has heretofore evolved largely on the heels of the various state data breach laws. Similarly, personal coverage is nascent, existing largely in the form of vendor-offered identity theft protection services; but if it were to gain ground, companies could arguably lose any incentive to lower their personal data emissions. The challenge with both personal and commercial liability insurance lies in the uncertain coverages and pricing due to the dynamic and interrelated nature of IoT privacy threats, not unlike the situation that exists currently with cyber insurance for security failures.

163 Tyler Moore, personal correspondence, December 2017.

164 See, e.g., Federal Trade Commission, 2017

NATURE OF THE IoT ECOSYSTEM

If personal data emissions can be narrowed to their *point sources* – emissions from identifiable and specific locations[165] – then market approaches are a better strategy, because they are more distinguishable and more controllable compared to other distributed approaches. Non-point sources of personal data emissions, which will likely be the case with IoT, are challenging to monitor and enforce and are therefore not conducive to regulatory mechanisms. When both data emission sources are present, a tax-based approach is more aligned when data handlers are identifiable, and subsidies or permits work better when data control is dispersed and unknown. Direct regulation is a favorable incentives approach where personal data is sticky and is retained by ecosystem actors.

As for incentives, subsidies and disclosure approaches transfer the costs and burden of proof compliance to regulated companies rather than the government, which can be appropriate when firms are in a better position to monitor and report their emissions. Tax or regulatory prohibitions are likely to fail because of costly enforcement and widespread noncompliance.

165 US Environmental Protection Agency, n.d.

UNCERTAIN COSTS AND BENEFITS

If the *cost of abating* personal emissions is uncertain and it is important to prevent polluters (e.g., large web platforms, consumer data brokers, credit reporting agencies) from bearing potentially high costs, then price-based market approaches are appealing. If there is more uncertainty with regard to the *benefits of controlling* personal data emissions and the objective is to prevent high environmental costs (negative externalities), then a quantity-based market instrument is more prudent (e.g., cap and trade).

One suggested course of action for incentivizing privacy risk management by both the supply and demand sides[166] is by granting legal property interests in personal data, such that individuals could be compensated when their data is used for commercial purposes.[167] Criticisms of this approach include: markets cannot function efficiently if property rights accrue to data; personal data is an extension of one's person and individuals should not be able to alienate a basic right; only the wealthy would benefit in a data-as-property regime; US property law would have to undergo a revolution to account for personal data; regulation would be needed because markets do not respond to individual privacy needs; there would be significant increased costs to businesses, including administrative costs and impediments to innovation that is predicated on the potential to benefit from data; and individual privacy rights will lead to market fraud.

Counterarguments, however, are: currently, individuals can already choose to alienate basic rights like freedom of speech in order to receive social benefits; disparate treatment toward the poor

166 The supply side are the manufacturers, developers, integrators, and retailers. The demand side are the users and entities seeking privacy protection.

167 Sholtz, 2001

is not likely since they would possess their identity as much as the wealthy; data property rights would simply be an extension of existing right to publicity (e.g., celebrities' property rights in photographs); businesses currently engage in a tremendous amount of 'wasted' advertising and marketing that produces no return on their investment, and consensual-based exchange of data would eliminate that guesswork; and there are non-economic benefits to privacy such as political freedom and liberty that comprise social democracy and welfare.[168]

Another potential source of incentives for companies to control personal data pollution lies with the investment supply chain. The investment community is an untapped mechanism that could potentially act as a forcing function to incentivize corporate social responsibility amidst a market that is lacking in incentives to embrace IoT privacy. Commenting on the positive influence that investors could have on IoT startups, Tien advocated:

I think it's really disingenuous in a way to hold up the college student innovators and say, "Oh, poor them," when you've got [these] experienced repeat players in the money business who have gone through the privacy issues with a LinkedIn or a Google or a Facebook, and already have a huge installed base and deep investment in those issues. It's not a new thing for them even if it's a new thing for the two guys who wrote an app.[169]

Implicit in Tien's statement is that investors are a natural choice for holding companies accountable for privacy pollution, because they have historical views into recurring privacy issues and can leverage their financial support to alter behavior. The reasons why investors should be interested in addressing privacy pollution is beyond the scope of this report.

168 Schwartz, 2004

169 Interview

In order to facilitate the supply side of IoT privacy management, it is essential to have a clear understanding of what constitutes 'sensitive data' and evokes a privacy risk. Additionally, companies interested in embracing privacy controls to reduce pollution need be able to grasp the benefits of maintaining lower data identifiability. The norms and standards that define privacy pollution need to go beyond conspicuous first-order identifiers to include the accrued privacy harm from derivative data fusion and linking (behavioral profiling, altered behavior, reduced autonomy). This will bolster the role of privacy as a factor in industry stakeholders' design and development models alongside considerations of cost, security, safety, accuracy, and reliability.

Finally, incentives for dealing with the disparate impacts of privacy management on innovation and market competition can include a combination of market instruments. These may be appropriate in helping correct market inefficiencies that result when privacy regulation disrupts companies that have market power, as in the health sector where innovation relies on privacy-sensitive data. If cost burdens were more equalized, some market instruments might favor incumbents, allowing them to control prices and creating barriers to entry. A permit system that sets aside a certain number of permits for new entrants might correct this problem.

In light of the current lack of market- and regulatory-based incentives for industry stakeholders to embrace IoT privacy management, we have proposed a path forward by modeling our recommendations for addressing personal data emissions by industry on the various incentive mechanisms currently used to address environmental pollution.

RISKS OF HARM

The IoT is the latest technical evolution to expose the divide between people's desire for privacy and the law's treatment of privacy harms. IoT privacy management, governance, incentives, and solutions are all bound by the notion of harm.

We use risk as a way to understand the likelihood that a harm will come to light. Our framing of IoT privacy harms in terms of risk reflects the fact that, in the technical environment of the IoT, some canonical harms[170] may not have manifested themselves yet, and some unforeseen harms may not yet be considered under current legal regimes.[171] Before describing those privacy harms, we first apply the principles of risk to the IoT environment in order to have a better grounding to understand the privacy harms that emerge in this new context.

Risk is typically defined as the severity and likelihood of occurrence of a harm to something of value (an asset) when a threat exploits a vulnerability. Our model for analyzing risk of IoT privacy harm frames the issue as a confluence of *threats* and *vulnerabilities* to privacy rights and interests – the *assets* in our application of risk.[172] This 'privacy risk equation' considers the potential for harm to both individuals and society as a result of threats that exploit IoT vulnerabilities. While this framing of risks of harm and its particular naming conventions may be unfamiliar to some of the privacy community, it mirrors core concepts from privacy frameworks such as the privacy impact assessment (PIA) - e.g., threats to personal data, potential risks, protection and mitigation strategies, controls, etc.[173]

170 E.g., economic loss from identity theft, fraud, loss of liberty from inaccurate information or improper use or exposure of information; and more traditional privacy torts: 1. Intrusion upon the plaintiff's seclusion or solitude, or into his private affairs. 2. Public disclosure of embarrassing private facts about the plaintiff. 3. Publicity which places the plaintiff in a false light in the public eye. 4. Appropriation, for the defendant's advantage, of the plaintiff's name or likeness. See Prosser, 1960

171 E.g., information collection by surveillance; aggregation of information; insecurity of information; and disclosure, exposure, distortion, and increased accessibility of information. See Solove, 2006.

172 See, e.g., NIST SP 800-30: http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf; The CNSS Instruction No. 4009: https://www.ecs.csus.edu/csc/iac/cnssi_4009.pdf

173 See, e.g., DHS PIA Template: https://www.dhs.gov/xlibrary/assets/privacy/privacy_pia_template.pdf; EU Privacy Impact Assessment Framework: http://www.vub.ac.be/LSTS/pub/Dehert/507.pdf

TECHNICAL UNDERPINNING

The technical characteristics of the IoT underpin and are the precursor for the privacy threats, vulnerabilities, and harms described subsequently. The architecture of the IoT reduces friction in collecting, processing, disclosure and actuation of data. The effect is a blurring of the temporal, spatial and organizational boundaries that have heretofore separated the physical, digital, biological, and social spheres. IoT sensors act as vectors for digitizing anything that can be sensed, capturing communications and visual, auditory, physical, and biological information that can then be managed, interconnected, and controlled.

The scope and degree of precision of IoT data capture moves the resulting information beyond fuzzy snapshots of human activities towards high-resolution data and inferences. The scale and opacity of data collection and flows will influence the relationship between individuals and organizations regarding the collection, use and disclosure of information, in ways that have important privacy implications. In sum, the technical drive to optimize and reduce friction in information flows results in increased friction for individuals attempting to maintain privacy through control of their information. We expand upon this further in the *Privacy Management* section.

PRIVACY THREATS

Privacy threats in the IoT are characterized by access, collection, use (analysis, actuation), and disclosure of personal data in violation of people's rights and expectations. As we discuss further in the *Boundaries* section, the scope of threats associated with the IoT is more expansive than the classic online realm, raising the probabilities of privacy risks. Broadly speaking, the threat landscape of IoT privacy should include the omnipresent attack vectors available to malicious actors.

We focus on the more tacit and understated portrayal of the threat model: the IoT drives equity conflicts between legitimate actors. In other words, the IoT may elicit a comparatively greater privacy threat by other, non-malicious ecosystem stakeholders – industry, the government, and other people – as a function of competing rights and interests that arise from IoT capabilities. Consider smart cities, for example, where sensors collect, analyze, and share data from light pole sensors that monitor vehicle and pedestrian traffic, parking, and local transportation: one person's expectation of privacy (to not be monitored or targeted) may conflict with the government's interest in enhanced public services, which may clash with a fellow citizen's expectation of safety, which may collide with industry's claim to commercial free speech (product and service offerings based on travel logistics).

These tensions represent another threat posed by the IoT in the form of a potential power imbalance. Governments are motivated to leverage personal information to secure infrastructure, interact with citizens, and serve other national interests. Industry is incentivized to exploit this same information to sell goods and services, and to protect the enterprise. In opposition to these

two actors, citizens have an interest in limiting data collection and ensuring their privacy, allowing data to be aggregated to promote anonymity, security, safety, and social democracy and regaining control of their economic interests. How these tensions resolve with regard to privacy in the IoT will be determined in part by the extent to which the IoT is leveraged to address power imbalances and associated information asymmetries unfavorable to consumers.

If one form of power is the ability to collect, process, and actuate data to exert control over individuals in ways that negatively impact their self-determination, the IoT threatens to shift the power landscape by exacerbating disproportionate control of personal information by companies and perpetuating a lack of the transparency which is so essential to consumers' accruing appropriate control. More equitable power – the chance for consumers to have a say in how their data is handled – would serve as a social and democratic check and balance. Power inequity, on the other hand, is a barrier to meaningful negotiations, competition and bargaining over rights and interests. In the IoT, power inequity will be a threat to privacy to the extent that data control is unchecked and consolidated by owners of platforms and services upon which consumers depend. If these IoT platforms are fueled by data from users, the users' lack of control over that data will threaten self-determination and ultimately create a self-perpetuating power imbalance.

The scope of the IoT threats is further complicated by a general lack of understanding of how threats will manifest. This makes it difficult to ascertain which precautions and mitigation measures to put in place to avoid or minimize adverse impacts. The Mirai botnet in the Fall of 2016 that commandeered hundreds of thousands of unwitting IoT devices to impose millions of dollars in damage from business interruption, fraud, and loss of data and customer loyalty[174] revealed the significant harm potential of using IoT devices to wreak financial, operational, and physical harm.[175]

174 Antonakakis, et al. 2017
175 Cogeco, 2017; Altman Vilandrie & Company, 2017

PRIVACY VULNERABILITIES

Another element in the privacy risk equation involves understanding the weaknesses or gaps in protections that can be exploited by the aforementioned threats to cause privacy harm. The scale and volume of data available for collection and use expands the range of opportunities to exploit data that implicates privacy and therefore increases the probability of realized harm. As discussed above in the technical underpinning, the digitization of anything that leaves a trace or is subject to sensing – biometrics, emotions, behaviors – introduces a privacy exposure point.

The traditional boundaries by which society has constructed privacy expectations are blurring, as we discuss in the *Boundaries* section. It is hard for individuals to know if the physical features that have assured a sense of solitude, permitted people to act anonymously, and supported control over identities are becoming ineffective. This is the case when data flows in an opaque, unobtrusive, automatic, regularized manner – all promised features of the IoT.

Even when users are aware of data flows, privacy vulnerability is increased by inadequate security of IoT devices. As well, even when deficient security is not the cause of privacy vulnerabilities, context-shifting and blurring between data collection for commercial and social settings creates another type of privacy vulnerability. Personal and social transactions and activities that are mediated by commercial information products and services are increasingly subject to commercial pressures to generate revenue. While the revenue model for the IoT is still emerging, it seems likely that many services will be

predicated upon users trading personal data for them. Consumers are being asked to provide and link more information to avail themselves of IoT functionality, yet so far have been given limited tools to control that personalization. Even when an individual is not the direct target of sensing, incidental data captured by other people's devices and the interconnectedness of large volumes of data increase privacy vulnerability.

In addition to the *degree* of vulnerability in the previous examples, the IoT introduces a relative difference in the *kind* of vulnerability that can enhance privacy risk. For example, sensing and digitizing sentiments and emotions yields a new path to measuring intimate features of people in ways not seen before.

To summarize, thus far we have applied the principles of risk to the context of IoT privacy to furnish a model to better understand, articulate and justify what constitutes privacy harms in the IoT that may not be sufficiently considered or foreseen by existing legal and technical controls.

IoT PRIVACY HARMS

Understanding the adverse impacts on individual and societal privacy rights and interests (the *assets*) comprises the third consideration in this framing of IoT privacy risk. We focus here on potential distinct differences in degree or kind of harms in the context of the IoT. Notably, this includes harms that the law has not codified but which the IoT evinces, such as social harms.

PERSONAL INFORMATION BREACHES AND IDENTITY THEFT

Conspicuous risks of harm in the IoT are those that ensue from inadequate security, such as breaches of personal information and identity theft.[176] If an IoT company loses data about users' personal behaviors gathered in their homes or in activities in public and their identity is linked to this data, this could cause measurable harm to consumers. Given the large numbers of data-collecting devices the IoT portends – a larger attack surface – breaches may increase. The resultant harms in such cases may be the most easy to quantify, relative to the other harms discussed below.

AUTONOMY HARM

The value of autonomy is inextricably bound up in the law's treatment of privacy harms. Indeed, transgressions of autonomy can be said to underline the majority of privacy harms, whether psychological (embarrassment, stigmatization, loss of trust, chilling effects on ordinary behavior, discrimination, intrusion on seclusion); economic (discrimination in employment, credit, education, and insurance) or physical (4th Amendment prohibition on unlawful search and seizure in the US). We consider anything that impedes people's self-determination while directly or indirectly engaging with information systems to constitute autonomy harm.

Collective autonomy harms can have disparate, far-reaching impact on the economic, physical, and psychological well-being of individuals and groups. Inadequate management of the personal data that informs the models produced by machine learning algorithms can result in public health and civil services disparities. Similarly, environmental sensor data can fail to aid vulnerable populations based on race or socioeconomic conditions. Collective autonomy harm can manifest as unequal access to and control of data. This risks engendering mistrust between individuals and institutions, resulting in impediments to or disengagement from social, political and economic activities that define individual and collective identities. Collective autonomy harms in the IoT warrant attention because if left unabated, these power imbalances get technologically embedded and institutionalized. They become hard to repeal and impact the entire fabric of social relationships within which privacy interests reside.

176 For example, in 2018 nearly 150 million users' personal details collected by the UnderArmour/MyFitnessPal app, including usernames, email addresses and passwords, were leaked in a data breach.

DIMINISHED USER PARTICIPATION

Both literature and our interviews point to diminished user participation as a potential harm in the IoT ecosystem. The 1980 OECD guidelines contained an *Individual Participation Principle*. It sought to give people the right to obtain personal data held about them by data collectors, to be given a reason if their requests were denied and an ability to challenge the denial, and to correct or erase incorrect or incomplete data. The US Fair Information Practice Principles also call for *Individual Participation* along nearly identical lines.[177] Further, the GDPR specifies a range of participation principles, including rights of access, rectification, erasure, data portability, and the ability to object to data processing and marketing.

Participation means having some say in the treatment of data about you. If privacy is "the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others,"[178] then Participation is an enactment of that claim. Naturally, one must know *who* has data about oneself to participate in its use, so Participation is necessarily reliant on transparency. Participation also requires the ability to *act* upon data about oneself – to affect the disposition of data, ergo the active verbs: access, rectify, erase, correct, object.

In part, Participation is diminished by a weakening of people's ability to knowingly consent to IoT data collection and use. A report by the US President's Council of Advisors on Science and Technology summarized the issue:

> *Cameras, sensors, and other observational or mobile technologies raise new privacy concerns. Individuals often do not knowingly consent to providing data. These devices naturally pull in data unrelated to their primary purpose. Their data collection is often invisible. Analysis technology (such as facial, scene, speech, and voice recognition technology) is improving rapidly. Mobile devices provide location information that might not be otherwise volunteered. The combination of data from those sources can yield privacy threatening information unbeknownst to the affected individuals.*[179]

177 US Dep't of Homeland Security, 2008
178 Westin, A., 1967, p. 7
179 President's Council of Advisors on Science and Technology, 2014

A concern over ambient data collection by the people near you also connects to consent – if a person doesn't know who is collecting information near them, there's no way to consent to that collection.

VIOLATION OF EXPECTATIONS OF PRIVACY

The law sets formal expectations of privacy rights. When there is incongruity between what society believes privacy harm to be and what can be remedied via the law, the result is fractured expectations of privacy. This incongruity in defining privacy harms emerges as a result of changes wrought by technology. Courts in the US are at least partially addressing the blurring between the concrete manifestations and abstract concepts of privacy injury when information is the proxy, but the general trend is slow to make the leap towards admitting these abstract concepts of privacy injury. Specifically, case law anchors on financial or physical harm that has provably already occurred and generally does not recognize the risk of future harm. Nor does it recognize negative impacts that are cumulative and collective.[180] For example, the threshold standard to bring about a lawsuit or achieve a favorable verdict demands 'concrete injury in fact.' Identity theft claims that rest on increased risk rather than actual fraudulent use are often deemed too speculative,[181] as are damage claims related to fear and emotional distress that result from increased vulnerability to a future attack. The difficulty calculating damages is one basis for dismissing such claims, as seen in June 2017 when a federal judge in California dismissed a class action lawsuit against Facebook for tracking its logged-out users' internet activity, in part because of the absence of sufficiently 'particularized and concrete' harm. [182]

Signals of fractured expectations of privacy are also evidenced in regulatory actions. However, at least as far as enforcement actions by the leading US consumer protection agency go, the gap between principles and implementation is smaller. The Federal Trade Commission has negotiated consent decrees based on non-monetary, abstract, autonomy

180 Solove, 2006

181 At the time of this writing, the 1st and 3rd US Circuit Courts of Appeals do not recognize; the 7th and 9th Circuits recognize an allegation of future harm if there is "danger of sustaining some direct injury" that is "both real and immediate"—such as identity theft.

182 Facebook Internet Tracking Litigation, US District Court, Northern District of California, No. 12-md-02314

and dignity-based harms.[183] In what is now an infamous, long-standing case, it declared 'substantial injury' to include intangible but concrete harm caused by the disclosure of sensitive medical information, declaring that "public disclosure of private information is by itself an actual concrete harm, even absent tangible effects or emotional injury" and, "a practice may be unfair if the magnitude of the potential injury is large, even if the likelihood of the injury occurring is low."[184]

ENCROACHMENT ON EMOTIONAL PRIVACY

As we discuss in the *Boundaries* section, the IoT is bringing devices ever closer to people's bodies. The increasing scale and proximity of sensors is beginning to open new frontiers in detecting people's emotional states and inner life. The potential resultant harm is an encroachment upon people's emotional privacy – a relatively new phenomenon being amplified by connected devices.

Prof. Andrew McStay cites experimentation with emotion detection in robots, in-home virtual assistants, children's toys, video games, and sex toys. Key reasons for the use of emotion detection are the ability to gauge reactions to content or advertising; affective behavior by a product, reacting to a child's or adult's emotions and dynamically changing the product's behavior to deepen and diversify its interactions; and nudging or manipulation in commercial contexts. "With insight into emotions, advertisers and retailers have a higher chance of influencing [behavior] and nudging us to spend,"[185] McStay writes. One of his interview subjects, representing a global trade association for the market research industry, expanded upon this, saying that:

> *The future of advertising and marketing lies in passive and always-on data collection, and that* ***the Holy Grail is real-time information about customer needs and emotions...*** *Today this is dependent on advances in mobile and wearable technology, and correlation of geo-location with contextual and behavioural information. The value of passive data collection is instant access to transactions and conversations... Seen this way, biosensors and biometric data promise additional real-time understanding as people move throughout everyday life, the city and retail spaces.*[186]

183 See, e.g., In the matter of DesignerWare, LLC, No. 112-3151 (Apr. 15, 2013).
184 LabMD, Inc. v. FTC, Case No. 16- 16270
185 McStay, 2018, Ch. 8
186 McStay, 2016, p. 4, emph. added

Another of McStay's respondents also related emerging devices to emotion detection. Dubious of traditional marketing surveys and focus groups, this Chief Insight Officer at an ad firm sought better methods to understand customers, "including biometric data about emotions to understand 'brand levers,' or how to get people to act, click, buy, investigate, feel or believe… [T]his objective involves 'understanding people through all of their devices and interaction points, i.e. wearable devices, mobile and IoT.' In general, the aim… is to 'collect it all' so to build more meaningful interaction with brands."[187]

McStay notes this executive's "interest in **emotional transparency**, or the unfolding of the body to reveal reactions, indications of emotions, feelings about brands, tracing of customer journeys and information that will help create 'meaningful brands.'"[188]

Given the advertising industry's growing hunger for the fullest range of human data, the incorporation of emotion detection across of a range of industries, the broadening capability of sensors, and penetration of those sensors deeper into private and public spaces, **it seems inevitable that the IoT will begin to challenge our emotional privacy**. The Internet of Things, bringing its sensors closer to bodies and thereby nearer the ability to, in Josh Cohen's terms, "externalize your inner self,"[189] brings to the fore a privacy risk that has been largely dormant. As to the harm, McStay succinctly asks, "if the body, emotions, and feelings aren't valuable in and of themselves, what is?"[190] There is a need for a social conversation about data, and the time to discuss the privacy impacts of emotion detection technology is *now*, while its use is not yet commonplace.

DIMINISHMENT OF PRIVATE SPACES

All of the privacy-challenging IoT characteristics we have discussed – proximity, scale, increased monitoring, boundary crossing, reduced ability to opt-out of collection – add up to an increasing diminishment of private spaces. This harms people's ability to achieve solitude and reserve, both from others and in their thoughts. We are concerned about a reduction in the availability of spaces for individuals to be able to retreat to and not be observed – spaces where one can

187 Ibid.
188 McStay, 2018, Ch. 8, emph. added
189 Cohen, 2014, p. 30
190 Interview

control who is present, who is listening, who is watching; places of seclusion. This diminishment of private spaces translates to a reduced ability to withhold data on lifestyle preferences, family dynamics, hobbies, etc. from third parties. A central site of this concern is the home. Heather Patterson of Intel commented that:

The bringing in of Alexas and Nest thermostats and DropCam cameras and Hue lights and everything being connected up, there's a lot of convenience associated with it, but I'm seeing in the research a lot of pervasive discomfort that's just arising from this feeling of being watched or never truly being alone.[191]

Meg Leta Jones worried about the impact of diminishing private spaces on children:

I think we should continue to find quiet, closed spaces that don't have any type of surveillance in them, and I think that kids should be really, really free from those spaces. That is part of them developing their own generation of an information society.[192]

Michelle De Mooy of the CDT put the diminishment of private spaces in the context of corporate actors:

You have companies that are enormous like Google and Facebook that control huge amounts of the economy, and it's arguable that their influence might be even more than the government's - they have access to vast amounts of personal data and therefore surveillance by [these companies] could be more damaging to privacy and ultimately more destructive to democratic principals.[193]

Privacy law scholar Joel Reidenberg discussed how the proliferation of technology that can identify people in both public and non-public spaces harms a long-standing expectation of privacy by obscurity:

Society can no longer claim any expectation of anonymity in crowds... The ubiquity of image-capture devices in private hands also means that individuals lose an expectation of privacy in non-public places... ***what might have been a previously obscure, anonymous presence in a private place becomes an identified act.***[194]

191 Interview
192 Interview

193 Interview
194 Reidenberg, 2014, p. 149, emph. added

Elegantly driving home Prof. Reidenberg's point, the New York Times reported in March of 2018 that Madison Square Garden, New York's famed event space, had been surreptitiously using face-scanning technology for security. A senior vice president of the company who manages the technology for the venue remarked, "The days of having 40,000 to 60,000 people in the stadium and not knowing who they are, I think those days are going to disappear."[195]

CHILLING EFFECTS

Lack of control over the constant data collection can result in psychological and behavioral chilling effects contrary to consumers' intentions. These effects may manifest as a reluctance to engage or trepidation when encountering IoT devices. The psychological insecurity about unwanted interference and manipulation may further manifest in people's actions. Feelings of malaise, resignation, or helplessness are subjectively real but, people may be unable to articulate legally cognizable harm because they lack information about whether or how data is actually being used.

Heather Patterson noted that people act differently when they feel watched:

People do watch those things in their homes that they are really concerned will expose them to judgment, to exclusion and isolation, and to feeling shame and feeling dirty and feeling lazy, feeling less than other people in some ways.[196]

Similarly, Brookman and Hans write:

In order to remain a vibrant and innovative society, citizens need room for the expression of controversial — and occasionally wrong — ideas without worry that the ideas will be attributable to them in perpetuity. In a world where increasingly every action is monitored, stored, and analyzed, people have a substantial interest in finding some way to preserve a zone of personal privacy that cannot be observed by others.[197]

NORMALIZATION OF SURVEILLANCE AND MANIPULATION

In 2014, Vint Cerf remarked in a report imagining life in 2025 that "continuous monitoring is likely to be a powerful element in our lives."[198] Here again

195 Draper, 2018

196 Interview
197 Brookman and Hans, 2013, p. 4-5
198 Pew Research Center, 2014, p. 16

is the common theme that the IoT is an evolution of prior historical trends – in this case, technology further normalizing surveillance, retrenching the existence of a *surveillance society*. This is a loaded term, of course, and requires disentanglement from its ominous connotations. Surveillance studies scholar David Lyon writes:

> ***Surveillance is part of the way we run the world in the twenty-first century.*** *Conventionally, to speak of surveillance society is to invoke something sinister, smacking of dictators and totalitarianism. But the surveillance society is better thought of as the outcome of modern organizational trends in workplaces, businesses and the military than as a covert conspiracy.*[199]

It is in this light that Cerf's comments should be read. Alongside – or despite – the utopian predictions of the IoT's near-term contributions to social welfare, connected devices must be recognized as sense organs in the surveillant assemblage[200] of the digital age. As Lyon explains further, "getting surveillance into proper perspective as the outcome of bureaucratic organizational practices and the desire for efficiency, speed, control and coordination does not mean that all is well. All it means is that we have to be careful identifying the key issues and vigilant in calling attention to them."[201]

De Mooy noted the tension between the Internet of Things and the democratic values of a freedom from surveillance and from government interference:

> *With devices in our homes that are maybe surreptitiously recording us or collecting and sharing information that we're unaware of, or even just the fact that they're in our homes recording, is that a violation of the idea that we are free from surveillance? From government interference? That in our homes, it is a private space? I think once that idea is challenged, then you have questions of is it possible to ever find a private space, and how necessary is private space to freedom of thought? I would answer that it's very necessary. It's at the core.*[202]

Just as the relentless, granular tracking of online activity has been normalized, **the Internet of Things will enable and normalize preference and behavior tracking in the offline world**. Indeed,

199 Lyon, 2008, emph. added
200 Haggerty and Ericson, 2000
201 Lyon, 2008
202 Interview

the characteristics of the IoT force one to ask if the very notion of an offline world will exist in the near future.

The emotional privacy concerns detailed above are undergirded by two emerging trends: the datafication of people's emotions so as to make everyday emotional life *machine-readable*; and subjecting emotion data to *commodity logic*. McStay observed, "In terms of where this is going, it's about attention, emotion, intention, and ultimately it's about commodification."[203]

Or, as scholar Shoshanna Zuboff puts it, the machine-readable human life represents "the migration of everydayness as a commercialization strategy… [reorienting] the nature of the firm and its relation to populations."[204] In a 2013 paper about self-generated health data, Heather Patterson wrote:

> *Encouraging… rigorous, whole body quantification of the self not only habituates individuals to the concept and practice of scanning and cataloguing activity and consumption habits, it results in corporations holding vast treasure troves of highly personal health data about tens of thousands of users, health and wellness libraries with unprecedented and complete entries of incalculable value to business associates, employers, and insurance companies.*[205]

We asked her what was the problem with this:

> *I feel like there's something kind of fundamental [quality] that can get lost… when we live in a society where… in the eyes of [some] others [we] exist just to be stripmined and have our individual particles re-assembled and converted to cash.*[206]

203 Interview
204 Zuboff, 2015, p. 76
205 Patterson, 2013, p. 9
206 Interview

Julia Powles echoed Patterson, titling a 2015 article, "We are citizens, not mere physical masses of data for harvesting."[207] Relatedly, one other researcher warned about how the information people contribute to their own commodification could be used against them:

If you think about products like Alexa... where sensing is being used to increasingly commercialize a home, that is inherently dangerous. It takes a space that nominally is not commodified, and commodifies it. So effectively, anything you say now can be used against your own economic interests.[208]

When asked what is at stake with the collection of intimate information and the datafication of everyday emotional life, McStay pointed out that this data allows corporations to manipulate and control human behavior with the aim of increasing consumption:

In terms of what is at stake, at a minimum, nobody could disagree that ***there's a better than average chance of raising the capacity to manipulate human behavior, typically in a consumer setting.***[209]

Prof. Ryan Calo puts such concerns in the context of the 'mediated consumer,' raising an alarm about the nature and *timings* of manipulating behavior:

[F]irms can increasingly choose when to approach consumers*, rather than wait until the consumer has decided to enter a market context... In an age of constant 'screen time,' however, in which consumers carry or even wear devices that connect them to one or more companies,* ***an offer is always an algorithm away.*** *This trend of firms initiating the interaction with the consumer will only accelerate as our thermometers, appliances, glasses, watches, and other artifacts become networked into an 'Internet of Things.'*[210]

207 Powles, 2015
208 Interview
209 Interview, emph. added
210 Calo, 2013, p. 1002-1005, emph. added

Worryingly, Calo notes how the American consumer protection regime is ill-prepared to address the potential harms of manipulating human behavior, stating that "market manipulation has had only a modest impact on regulatory policy."[211] This concern was brought into stark relief in 2017 and the beginning of 2018 when it was shown that Russian disinformation campaigns were conducted through Facebook to affect US voting behavior,[212] and that extensive profile data on over 87 million users was inappropriately collected by Cambridge Analytica, a political consulting firm, to use in their psychological profiling and targeting of US voters in the 2016 election.[213]

211 Ibid., p. 1002
212 Karpf, 2017
213 Chang, 2018

CONCLUSION

The IoT is an evolutionary step in the development of industrial, enterprise, and consumer technologies. The scale and proximity of sensor technology, ubiquity, opaqueness, and unobtrusiveness amplify power imbalances between citizens and data collectors, making them more vulnerable to a range of privacy harms. Data breaches and the loss of personal information will likely increase as the IoT creates a larger attack surface. People's expectations about the private nature of their activities, interactions, and behaviors are likely to be further challenged. As the human environment acquires more monitoring devices, private spaces into which people can retreat and find reserve may diminish. And, if people feel they are being watched, they might chill their speech and conform their behavior in response. The closeness of sensors to people's bodies implicates emotional privacy, especially considering there is a commercial interest in consumers' emotions as a way to better understand and sell to them. Broadly, the IoT portends a greater ability to render society machine-readable, enhancing surveillance and commodification. The tracking of behavior and preferences that are now a standard part of living online will migrate to the offline world. However, all is not lost – the next two sections, *Emerging Frameworks and Strategies* and *Recommendations*, discuss a range of ways to address the above risks and harms.

EMERGING FRAMEWORKS AND STRATEGIES

This section assembles the principles, practices, techniques, and frameworks that emerged from our interviews, workshops and literature to address the privacy risks examined in prior sections. They fall into three categories: *user control and management*, *notification*, and *governance*.

USER CONTROL AND MANAGEMENT STRATEGIES

To enhance people's control and abilities to manage personal data we find that there are pre-collection and post-collection strategies for organizations to implement, and identity management (IDM) strategies, which we separate because they cut across both pre- and post-collection.

PRE-COLLECTION

Data minimization

One strategy to improve the privacy posture of IoT devices is to not collect certain information at all. As the FTC has stated: "Thieves cannot steal data that has been deleted after serving its purpose; nor can thieves steal data that was not collected in the first place."[214] Avoiding data collection not only reduces the possibility of data loss, but it may reduce cost or computational overhead if that data is truly not needed for features to function. A privacy impact assessment (discussed below) can surface which data elements could be eliminated. Some data protection laws use the term 'purpose limitation' to refer to the practice of restraint in data pre-collection strategies: if the data does not serve an immediate business or essential functional purpose, do not collect it.[215]

214 FTC, 2015, p. 35

215 See Art29WP Opinion 03/2013 on Purpose Limitation: http://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2013/wp203_en.pdf

An illustrative case of preventable automatic data collection is vehicle infotainment systems. Plugging a phone into a car's USB port can sometimes trigger an automatic data transfer, pulling across people's contacts, GPS information, phone number, and other personal data, often without appropriate or granular consent.[216] Many infotainment systems do not provide obvious or user-friendly ways to delete such data. In the case of rental cars, this raises obvious privacy concerns, as it leaves a very detailed set of personal information lying unprotected in a quasi-public space. In 2016, the FTC issued a warning on this issue, suggesting car USB ports should not be used at all for charging.[217] The design imperative here is clear: always ask for data-specific permission before downloading phone information, do not collect irrelevant information, and make it easy for people to delete all personal data. For example, in addition to limiting the data gathered by infotainment system, a privacy-by-design approach would suggest that rental car companies should have a standard policy to wipe all personal data from cars right after they are returned.

Do Not Collect 'Switches'. *Devices can and often should be designed with easy ways to shut down collection sensors. Research on RFID privacy in the early 2000s suggested that users have a right to disconnect from their networked environment, and so should be able to easily deactivate the tracking functions of RFIDs. French Internet expert Bernard Benhamou coined the term the 'silence of the chips' to capture this notion,[218] and the Art29WP[219] has recently updated the idea, stating: "Similarly to the 'Do Not Disturb' feature on smartphones, IoT devices should offer a 'Do Not Collect' option to schedule or quickly disable sensors." This could be interpreted to mean actual physical switches, or obvious toggles within an interface.*

216 Akalu et al., 2016
217 Schifferle, 2016
218 Santucci, 2013
219 Art29WP, 2014, p. 22

Wake Words and Manual Activation

While a common defining characteristic of the IoT is that devices are always-on, from a design perspective there's much more nuance to consider. Stacey Gray of the Future of Privacy Forum helpfully illustrates the granularity of 'always-on,' proposing three categories of microphone-enabled devices: manually-activated (e.g., a switch or button), speech-activated (via a 'wake word'), and truly always-on (constantly transmitting data).[220] The strongest privacy choice is one where users are aware that a device is listening or watching, such as via a red light or other indicator, and must take a clear action to activate it. In line with the 'Do Not Collect switches' described above, Gray notes that a hard 'mute' button on devices with wake words is a privacy-enhancing choice, curbing unintended audio capture and alleviating concerns about surveillance and intrusion.[221]

Privacy Impact Assessments

One way to assist all the above design considerations is by performing a privacy impact assessment (PIA), what the GDPR calls a Data Protection Impact Assessment (DPIA). PIAs are systematic processes to evaluate the impact and risks of collecting, using, and disseminating personally identifiable information in a project, product, service, or system. The goals are to identify privacy risks; ensure compliance with national or local laws, contractual requirements, or company policy; and put risk mitigation strategies in place. PIAs help organizations get a better sense of the personal data they handle, as well as understand the associated risks and how to manage issues in case of a data breach or accident.[222] Basic elements comprising a PIA include:

220 Gray, 2016
221 Ibid., p. 9

222 See Wright and DeHert, 2011

- Data sources
- Data flows through the product/service lifecycle
- Data quality management plan
- Data use purpose
- Data access inventory—who inside and outside the organization can access the data
- Data storage locations
- Data retention length
- Applicable privacy laws, regulations, and principles
- Identification of privacy risks to users and the organizations and the severity level (e.g., High, Medium, Low)
- Privacy breach incident response strategy

In the US, PIAs have been mandatory for federal agencies for some time, but they remain uncommon in industry. The GDPR requires data controllers[223] to perform DPIAs when data processing is "likely to result in a high risk to the rights and freedoms of natural persons."[224] The Art29WP released guidance on how to interpret this requirement. The list of activities that would trigger a DPIA is somewhat complex, but the following are particularly germane to the IoT:

- *Systematic monitoring*
- *Sensitive data or data of a highly personal nature*
- *Innovative use or applying new technological or organizational solutions*[225]

223 In GDPR terms, the primary entity that collects, stores, and directs the processing of personal data. External entities that process data on behalf of the data controller are called 'data processors.'
224 GDPR Art. 35
225 Art29WP, 2017, p. 9-10

The GDPR applies to all companies who process Europeans' personal data, including American and other countries' companies who serve customers in the European Union, whether or not they have a physical or legal presence in a European member state. For the first time, the GDPR puts in place a sanction regime for companies who do not carry out an assessment when they should: up to 2% of annual turnover or €10 million, whichever is greater.[226]

POST-COLLECTION

Data Deletion

Continuing from *Data Minimization* above, deleting data is a very good risk-reduction strategy. This could mean:

- deleting data as close to the sensor as possible
- deleting data after aggregating it
- deleting data after a period of time has elapsed and the data has become 'stale'

The GDPR creates a *right to erasure*. As with many aspects of the GDPR, this right is not straightforward and has counterbalancing aspects, but in general European Union data subjects have the right to cause data controllers to delete data about them if they withdraw consent or when the data is no longer needed for the purpose for which it was collected or processed.[227] There is no substantive regulatory equivalent in the US,[228] but it remains a principle that enhances a user's participation rights.[229] Giving users the ability to easily delete data supports their ability to make substantive choices about their personal information.

226 Ibid.

227 See Maldoff, 2016 for a brief overview

228 The exception to this is if a company promises in their privacy policy that they will give users the ability to delete their data, they must fulfill this promise else be exposed to FTC action for deceptive practices.

229 See Risks of Harm section

Consent withdrawal

Requiring consent before data is collected or processed is a key feature of privacy and data protection regimes, but withdrawing that consent is problematic. Arguably, consent is meaningless without an ability to withdraw it, to say nothing of granular consent possibilities ('yes for *this* purpose, no for *that* purpose,' etc.). The GDPR allows for the possibility of preventing further processing by mandating that "it shall be as easy to withdraw consent as to give it."[230] Such withdrawal would then be a prelude to triggering one's right of erasure. Ease of consent withdrawal and deletion are both supportive strategies for user choice, control, and participation.

Encryption

Traditional encryption schemes were designed with greater resources in mind. As a result, many of them will not work well with IoT devices, which challenge the basic strategy of encrypted data at rest and in transport because they are 'resource-constrained': low power, small form factor, low memory, low processing power, low cost, limited network bandwidth, low storage capacity. This has spurred much research and development on lightweight cryptography, which takes IoT devices' constraints into consideration. NIST's 2017 report on the subject is a useful starting point for exploring lightweight encryption as an IoT privacy and security strategy.[231]

IDENTITY MANAGEMENT

Since the early 2000s, the identity management (IDM) community has actively explored many modern privacy issues: the separation of informational contexts, cross-correlation of online activities, the value of pseudonymity, the sensitivity of identification, uncontrolled profiling, and the need for personal data management systems to be designed for and controlled by users.[232]

230 GDPR Art. 7(3)

231 McKay et al., 2017

232 See, e.g., ICPP and SNG, 2003, "Identity Management Systems: Identification and Comparison Study": https://www.genghinieassociati.it/wp-content/uploads/2011/11/ICPP_SNG_IMS-Study_Summary.pdf

IDM research and commercial work has yielded valuable language and concepts, such as: *unlinkability* – the intentional separation of data events and their sources, breaking the 'links' that form between the different places users go online or between different devices; and *unobservability* – building systems that are blind to user activity. Both of these can be employed in the design of IoT devices and their host systems. For example, unlinkability in an IoT context can be restated as: my car/fitbit/voice-assistant (and therefore, the manufacturers and intermediaries) do not need to know which websites I visit, or which other devices I use. Unobservability is a design principle that can be applied to intermediaries and IoT platforms, blinding them from user activity. These two concepts are a form of data minimization, limiting which parties can see personal data.

The identity dimensions of the IoT are just beginning to take shape, and there are two key considerations: machine identity and human identity. Machine identity concerns the ways that devices are authenticated and kept track of within an IoT management system. This encompasses security, machine 'trust,' and the provenance of the data coming off a device. Human identity, in the IoT context, concerns questions which all have privacy implications:

- Who is the device owner?
- Are there additional users?
- Is there an option for unidentified Guest Users?
- Is there an option for pseudonymous use?
- How are users' system rights managed?
- Can users be given delegated/partial control rights? (E.g., change the heater temperature but not turn it off)
- Who can access data stored on a device or direct sensor feeds?
- Can device owners see other users' data?
- Can two users see one another's data or change each other's settings?
- Can data from a device be verifiably associated with a particular user?
- How do users disconnect/disassociate themselves?

Identity management is a useful lens through which to view IoT privacy as it moves the discussion away from intrusion, unexpected practices, breaching confidentiality, and threats to seclusion into a discussion of privacy as control, access management, and selective sharing. That is, IDM concerns itself with *architectures of permission*. IDM expert Eve Maler observed:

> *IoT data is available for you to share, but you may not wish to share it with everybody in the world. It's selective sharing. Privacy has historically meant, "Stay away from me and don't take my data," but it's starting to appear to me in a model where if it's user-controlled, person/individual-controlled, data-subject controlled, then 'controllable' means there's an ebb and flow. 'Controllable' means you get to release it when you want to.*[233]

Maler uses the example of the Google Docs Share button as a model for delegated access and selective sharing of both data and system functions. The Share button allows a document owner to grant others the ability to edit, comment, or only view its contents. Maler asks, "Is that opt-in? Is that opt-out? No. But it's permissioning of a very powerful sort."[234]

In her work, Maler is seeking to "shift the consent concept to permissioning."[235] Maler illustrates what she calls "relationship-driven permissions" through the example of a family home and an Airbnb rental. The home contains the home owner and family members; it has many 'smart' devices in it, and each resident can have different permission rights. When the family rents out their entire home on Airbnb, the guests are given a narrow set of rights to use the smart lock, the air conditioner, the speakers, internet access, smart lighting, and so on. When the guests leave, their access and other rights are revoked.

As an example of work which aims to change the non-granular, binary conception of consent to multi-modal, granular permissioning arrangements, Maler points to the emergence of 'consent tech.' These are particular families of standards and technology that build upon each other to yield frameworks that address different pieces of Maler's vision of control and selective sharing:

233 Interview
234 Maler, 2017
235 Ibid.

- OAuth 2.0 – a protocol for authorization flows for applications and devices
- OpenID Connect – an identity authentication layer built upon OAuth 2.0
- User-Managed Access (UMA) – a protocol designed to give people a unified control point for authorizing who and what can get access to their digital data, content, and services, irrespective of their location
- Health Relationship Trust (HEART) – patient-centric health data sharing from heterogeneous sources; based on the three above technologies[236]
- Consent receipts – track a user's consent, linking it to a privacy policy, making it available across organizational boundaries[237]

According to Maler, these technology families help systems meet high privacy implementation standards such as those required in the GDPR for unambiguous consent, proof of that consent, data minimization, accuracy, and quality.[238] IDM technology also directly addresses boundary management concerns by providing tools to control how we let people and organizations in and keep them out. More broadly, all of the above identity management technologies, concepts, and perspectives are helpful discussion topics when attempting to envision a world where people have fine-grained control over the sharing of data collected by their devices. If the IoT does indeed herald a much more monitored environment, then strategies and technologies that help people shape their informational life through selective sharing are essential.

236 See http://openid.net/wg/heart/
237 See https://kantarainitiative.org/confluence/display/infosharing/Consent+Receipt+Specification
238 Maler, 2017

NOTIFICATION STRATEGIES

Both the US and Europe are heavily invested in Notification as a privacy and data protection strategy. While there is vigorous debate as to its effectiveness,[239] it remains a bulwark of modern privacy regimes. Fortunately, there is active research and policy discussion in the following areas:

NOTICE TIMING

Research has shown that the *timing* of notices has an impact on their effectiveness. Notifications are most commonly displayed to users at setup, when programs, services, or devices are first used. But notifying users at the time when data is being actively collected or used is also a good way to get a user's attention. These 'just-in-time' notices support people's ability to make a decision about whether to accept or reject a data practice. Crucially:

> *Just-in-time notices and obtaining express consent are particularly relevant for data practices considered sensitive or unexpected. For instance, in the case of mobile apps, access to sensitive information such as the user's location, contacts, photos, calendars, or the ability to record audio and video should be accompanied by just-in-time notices.*[240]

239 See, e.g., Calo, 2012; Sloan and Warner, 2013

240 Schaub et al., 2015, pp. 6-7

The California State Attorney General,[241] UK Information Commissioner's Office[242] and Art29WP[243] have all endorsed this strategy and developers have been implementing it for some time. A common example is when a browser or mobile app asks for permission to know your location or access your camera or contacts.

Researchers also endorse three other timing-oriented strategies: periodic, context-dependent, and persistent.[244] *Periodic notifications* remind users of ongoing data practices, giving them the opportunity to reaffirm or adjust their consent over time. Schaub et al. write, "Periodic reminders of data practices can further help users maintain awareness of privacy-sensitive information flows. **Reminders are especially appropriate if data practices are largely invisible."**[245]

Example of periodic notification

Context-dependent notices appear based on changes in a user's or system's context, such as when a user enters or leaves certain locations, or, in the case of social media, changes the intended audience of a post. In theory, an IoT device could display privacy choices when it moves from inside the home to outside it, or from when a user is alone to when there are others present. An advanced form of this notice could inform the user about potential inferences that could be made based on data gathered in differing contexts.[246]

Persistent notices provide awareness of ongoing data practices, such as icons in the status bar of mobile phones to indicate that location data is being used, or indicator lights when a device is recording. Since users can become quickly inured to their presence, though, a "system should only provide a small set of persistent indicators to indicate activity of especially critical data practices."[247] Nor is it efficient or effective to display everything about a system's data practices in one notice at one time. Instead, privacy notices should be *layered*: "notices shown at different times, using different modalities and interfaces, and varying in terms of content and granularity in a structured approach."[248]

241 California Dep't of Justice, 2013
242 UK Information Commissioner's Office, n.d.
243 Art29WP, 2017, p. 28
244 Schaub et al., 2015
245 Ibid., p. 7, emph added
246 Schaub, Interview
247 Schaub et al., 2015, p. 7
248 Ibid. p. 5

NOTICE COMPREHENSION

It's critical, of course, that notices be *understood* by their audiences. The opacity and density of most existing privacy policies are central complaints and hindrances; they are 'written by lawyers for lawyers.' Each of the strategies listed above attempts to improve comprehension in one way or another, but the question remains: how can companies be encouraged to make their privacy policies understandable and useful? A key answer is that they must *test* their notices for comprehensibility:

> *There needs to be testing that shows that this notice is actually effective and first of all, being noticed. That people see that there's a notice, that they can understand the notice, and that they can act on the notice.*[249]

Florian Schaub further noted that, within industry, there is often a fear that moving away from *one* defined privacy policy to something more integrated into the user experience could create duplicate messages that might not provide the comprehensive picture of a data practice that a privacy policy would. This may be perceived as a legal or compliance problem. The solution to this is for regulators to release guidance on the type of notice that is permissible and complies with existing regulations, and what level of innovation or leeway is acceptable.[250]

In response to these issues, the FTC held a workshop on disclosure evaluation in September of 2016. This important, far-ranging workshop explored comprehension, attention, cognitive models, measurement, decision-making, and lowering the cost of notice testing. The resultant Staff Report[251] is an essential primer for both industry and regulators.

One burgeoning area of notification research is the attempt to *automate* different parts of privacy notification, preference determination, and choice. Interdisciplinary researchers have been exploring ways to allow software to learn about users' privacy preferences by observing their choices and behavior over time.[252] Based on what these software agents learn, they could automatically configure a user's privacy settings within a mobile phone or IoT device, or offer suggestions. These agents could also periodically 'nudge' users to examine or

249 Schaub, Interview

250 Ibid.
251 Federal Trade Commission, 2016
252 Sadeh, 2017

revisit privacy options. Early research shows that these nudges are effective in motivating users to engage more with their settings.[253] Automation is also being employed to help users be more aware of IoT devices in their environment. Researchers from Carnegie Mellon University have developed a 'privacy-aware notification infrastructure,' where IoT devices broadcast their existence and their sensing capabilities, which are then collected by an IoT resource registry, and users are notified by an IoT Assistant on their mobile device.[254] Such a design engages directly with the privacy problems of device invisibility and their lack of substantive interfaces.

NOTICE CONTENT

As we discussed in the *Privacy Management* section, there is a troubling pattern emerging of manufacturers not sufficiently disclosing how devices collect data. Regulators should issue guidance on privacy policy disclosures about the nature and use of sensor data: what sensors are on the device, what data is gathered, where it is stored, and whether it is encrypted or de-identified.[255] The same is true of the types of inferences that can be derived from sensor data and the types of analysis applied to the datasets, without necessarily delving into trade secrets. Further, when claiming that personal data has been de-identified, the standards used to do so should be disclosed.

NOTICES CONTRIBUTING TO NORMS

Finally, IoT privacy policies should commit companies to the principle that consumers own the sensor data generated by their bodies, cars, homes, phones and other connected devices.[256] While creating proprietary interests in personal data is an extremely complicated proposal,[257] the *norm* created by such a commitment would be highly advantageous, amplifying citizen's expectations about the value of their personal data and how it should be stewarded.

Taken together, the above notification enhancement strategies represent some of the best research and development available to increase people's awareness of how data about them is gathered and used. The strategies provide meaningful ways for people to affect the data practices going on around them, while also reducing the cognitive burdens involved in doing so.

253 Ibid.
254 Ibid.
255 Peppet, 2014, p. 162
256 Ibid.
257 Prins, 2006

GOVERNANCE STRATEGIES

This section covers legislative strategies, as well as public-private and market-driven governance strategies. We discuss emerging law, suggested legal changes and rights, risk standards and certification regimes.

EU GENERAL DATA PROTECTION REGULATION

The most significant governance development pertinent to the IoT is the EU's GDPR. While we have highlighted specific elements of the GDPR in prior sections (such as the PIA section above), a comprehensive review of this expansive law is beyond the scope of this report, but it will clearly affect how companies collect data and notify users about its usage. These are early days for the law and there is sure to be much ambiguity in the months and years ahead. Still, with its global focus and expensive sanctioning power, the GDPR's impact on the privacy posture of devices and the services behind them will potentially have wide effect. Europe's ePrivacy Directive, which is still being negotiated, is expected to play a similarly strong role. It specifically cites the IoT, noting that the principle of confidentiality applies to machine-to-machine communications, and that additional specific safeguard could be adopted under sectoral regulation.[258]

258 European Commission, 2017b, Recital 12

BASELINE PRIVACY LAW IN THE US

As we mentioned in the main *Governance* section, one oft-discussed strategy is the enactment of baseline, omnibus federal privacy law for the United States. This would help to clarify privacy expectations for the populace, and improve the state of privacy that results from a porous, sectoral regulatory regime. The European model need not be absorbed wholesale; it serves as an instructive example. If indeed sectors and social contexts are collapsing, an omnibus, baseline federal privacy law is one method to strengthen privacy as the number and diversity of IoT devices increase.

DELETION RIGHTS

The GDPR stipulates a right to erasure in cases when data is no longer needed for processing, when people withdraw consent, and other reasons.[259] It's not an absolute right, and it may prove difficult to exercise in some cases, but it establishes the core principle that people have a right to cause organizations to delete data about them. This is a new and important privacy right, speaking directly to issues of autonomy, choice, and strong user control. Sometimes called 'the right to be forgotten,' this right has been the source of much contention. Detractors see it as a step towards erasing historical memory or enabling the hiding of past misdeeds. The GDPR goes out of its way to address such criticisms, carving out exceptions for public interest, public health, and scientific or historical inquiry.[260] Still, the power of this new right should not be overlooked, and non-European countries should also implement it in service of progressive policy. As the IoT is likely to collect enormous amounts of personal and sensitive data, and since people may not be aware of the scope of it, a right to erase data is a valuable way to give people enhanced control over that data. California has implemented a limited form of deletion rights for

259 GDPR, Art. 17
260 Ibid.

children, requiring websites, apps, and other service providers to make their content or information *invisible* upon request (but it may stay resident on the providers' servers).[261]

The chance that deletion rights will emerge in US national legislation is very low. That said, the norm is important – it creates an expectation that organizations should delete data when a person wishes them to. Businesses often grant rights, powers, options and protections that the law does not mandate, such as the periodic notification examples we cite above. There is nothing preventing US businesses from offering people the right to have their data deleted when they stop using a service, and they should be urged to do so by the privacy community and the public. Deletion rights are especially valuable for people whose data was collected when they were children, but we argue that the principle has general application. It supports the view that people are allowed to change their minds about consent and data sharing, and it can help rectify data collection users were unaware of.

USE REGULATIONS

Given the view that Notice and Consent may be failing strategies in the face of ever-increasing data collection, some have suggested regulating data by *use*, rather than trying to declare all expected data uses up front at the time of consent.[262] In other words, after collection, some data uses would simply be disallowed. One of the world's first data protection laws, the US Fair Credit Reporting Act, is a use regulation: it specifies that consumer credit data may only be used in creditworthiness, employment, and housing decisions.[263] With regard to the IoT, we've identified the problems with data collected in one context leaking into others. Not only can this violate users' expectations, but there is the danger of discrimination when unwanted parties learn about personal characteristics, such as the case of an employer learning about an employee's personal or political activities. Prof. Scott Peppet argues that privacy advocates and consumer groups should focus on keeping IoT data from violating contextual boundaries. For example, IoT health and fitness data or details learned from in-home devices should be restricted from use by insurers, lenders and employers:

261 CA SB-568

262 Cate, et al., 2014
263 Fair Credit Reporting Act, 15 U.S.C. § 1681

A woman tracking her fertility should not fear that a potential employer could access such information and deny her employment; a senior employee monitoring his fitness regime should not worry that his irregular heart rate or lack of exercise will lead to demotion or termination; a potential homeowner seeking a new mortgage should not be concerned that in order to apply for a loan she will have to reveal her fitness data to a bank as an indicator of character, diligence, or personality.[264]

Similarly, people should neither be forced nor economically coerced into revealing data collected by IoT devices. Peppet notes:

One can easily imagine health and life insurers demanding or seeking access to fitness and health sensor data, or home insurers demanding access to home-monitoring system data. As such data become more detailed, sensitive, and revealing, states might consider prohibiting insurers from conditioning coverage on their revelation.[265]

Such types of prohibitions are already being enacted to prevent car insurers from conditioning the sale or claim payment of an insurance policy upon gaining access to a vehicle's event data recorder ('black box').[266] A central point of contention in the use regulations debate is whether it should be businesses or regulators that specify which uses violate expectations.[267] The kinds of prohibitions Peppet suggests would be legislatively driven, but there is also an incentive for businesses to consider restricting their data uses in service of consumer acceptance of IoT technologies. For example, businesses have wide latitude in which *defaults* they choose to enable in their products. Starting from the idea that certain data is sensitive, such as health and fitness information, manufacturers can choose to deploy devices with sharing and publicizing turned off. This comports with the Art29WP's view: "Information published by IoT devices on social platforms should, by default, not become public or be indexed by search engines."[268]

264 Peppet, 2014, p. 151
265 Ibid., p. 151

266 See Rosner, 2016a, pp. 3-5; Peppet, 2014, pp. 152-155
267 See Rosner 2016b
268 Art29WP, 2014, p. 23

RISK STANDARDS

Standards that address risk management in the IoT are nascent at best. There is opportunity to leverage guidance from other domains whose institutionalized risk standards have proven to be capable. Examples include:

- US Dep't of Health and Human Services Risk of Harm Threshold for Breach Notification under the 'HITECH' Act: a breach is a use or disclosure that "compromises the security or privacy of the protected health information" that "poses a significant risk of financial, reputational, or other harm to the individual." A breach is presumed to have occurred unless the business proves that there is a low probability that protected health information has been compromised.
- FTC's standard for Fairness: [1] the act or practice causes or is likely to cause substantial injury to consumers [2] which is not reasonably avoidable by consumers themselves and [3] not outweighed by countervailing benefits to consumers or to competition.[269]
- De-Identification: The de-identification (de-ID) of data is maturely practiced in the health field, and there are several methodologies that facilitate the management of re-identification risk. NIST's Internal Report 8053[270] and the HITRUST De-Identification Framework[271] are valuable collections of de-ID methods and frameworks. The Future of Privacy Forum has helpfully assembled "A Visual Guide to Practical De-Identification" that displays the spectrum of fully identifiable personal data to aggregated anonymous data.[272]

269 FTC, 1980
270 Garfinkel, 2015
271 HITRUST, n.d.
272 Finch, 2016

SECTOR-SPECIFIC PROTECTIONS: WEARABLES

In their comprehensive report, "Health Wearable Devices in the Big Data Era," the authors lay out a set of actionable requirements for the health and fitness wearable sector:

- *All* data collected from a health or wellness wearable device should be considered sensitive, and thus require an affirmative and effective consent process before they can be collected and used.

- Clear, enforceable standards should be established for both the collection and use of information on wearables and other Internet-connected devices, with allowances for consumers to place limits on the data collected by and about them.

- Companies should be required to explain fully and in clear language what their data practices are, and there should be standardization of terminology so that comparisons are possible. They should also be required to make public disclosures about how they analyze data and use the derived knowledge.

- Wearable and other connected-health companies should not share user information with any third parties where advertising, marketing, or the promotion of other services are involved.

- Companies should comply with a person's request for her own data as soon as possible and at the lowest cost.

- The metrics used to determine how de-identification is most effectively accomplished should be disclosed and subject to independent verification.

- Wearable devices and apps should be tested to determine that consumers will be able to understand their privacy choices and terms of services.

- Self-regulatory organizations should develop standards that apply to all sectors of the consumer connected-health industry, along with a process for independent auditing.

- The various participants in the digital health sector, including the wearable and mobile apps industry, should develop a set of fair marketing practices for using health-related data.[273]

CERTIFICATIONS

Various government and private sector bodies have begun publicly discussing what an IoT certification scheme might look like. In 2016, the European Commission (EC) expressed support for a Trusted IoT Label as part of its larger Digital Single Market program to help consumers better understand varying levels of privacy and security in IoT products. The EC likened this idea to energy efficiency labeling requirements, in place since 2010.[274] In September of 2017, the EC released a proposal for new regulations to bolster cybersecurity in the European Union, in which they discuss the role of cybersecurity certification in relation to the IoT: "The digital single market, and particularly the data economy and the Internet of Things, can only thrive if there is general public trust that such products and services provide a certain level of cybersecurity assurance."[275]

273 Montgomery et al., 2017, pp. 6-7

274 European Commission, 2016a, footnote 88
275 European Commission, 2017a, p. 31

The private sector has also begun to propose certification schemes. The UK nonprofit IoT Security Foundation has launched a voluntary, self-assessed Best Practice User Mark "intended to help users communicate publically that they take IoT security seriously, are IoT security aware and are conscious of their responsibilities as a supplier of IoT products or services."[276] Also in the UK, work on an 'IoT Mark' has begun that would be formalized as a consumer-facing certification mark under British law. Growing out of the London IoT Meetup community, this mark scheme has attracted the attention of a wide range of stakeholders and will cover[277]:

- Data security
- Customer and consumer privacy
- Data governance
- Hardware & software security
- Interoperability
- Provenance
- Lifecycle
- Accountability in the supply chain

276 IoT Security Foundation, n.d.
277 See https://iotmark.wordpress.com/

In the US, Consumer Reports has partnered with Ranking Digital Rights, a nonprofit research project that reviews commercial privacy policies; and the Cyber Independent Testing Lab, a nonprofit software security testing organization founded by security expert Peiter 'Mudge' Zatko and others, to create The Digital Standard: "an ambitious, open, and collaborative effort to create a digital privacy and security standard to help guide the future design of consumer software, digital platforms and services, and Internet-connected products."[278] The Standard uses a wide-ranging set of 'tests' to evaluate IoT products and companies for characteristics such as safety, product stability, required password strengths, security lifecycle, willingness to disclose vulnerabilities, degree of data control by users, data retention and deletion, data minimization, transparency, and corporate social responsibility.[279]

The schemes listed above are potential components of IoT governance, norm setting, consumer purchasing behavior, government procurement, and product liability. The Digital Standard, in particular, is the most comprehensive attempt to unify product features with value-laden characteristics, such as commitment to human rights, under a single certification banner.

278 See https://www.thedigitalstandard.org/
279 See https://www.thedigitalstandard.org/the-standard

RECOMMENDATIONS

In this section we make recommendations to improve the state of IoT privacy based on areas we find to be either unaddressed or weakly addressed by current efforts, as well as areas where existing strategies should be amplified. The recommendations fall into four categories: *research*, *funding*, *fostering discussion*, and *governance*.

RESEARCH

EMOTION DETECTION

Much more research is needed on the subject of emotion detection, in particular:

- People's expectations about the gathering and use of emotion data, both from adults and children

- Policy weaknesses around the collection of emotion data
  - > US: how it fits within the PII framework, and the adequacy of this approach; whether the Children's Online Privacy Protection Act (COPPA) is adequate for the collection of children's emotional data
  - > EU: the evolving interpretation of emotion data under the GDPR
  - > How policy and the law deal with 'intimate' but not yet identifiable personal data
  - > Consumer protection dimensions of stockpiles of emotion data

- Security and privacy of existing emotion detection products (analyses should be performed by researchers and commercial firms)

CHILDREN'S ISSUES

Most privacy research is focused on adults. Given that there will be connected toys marketed specifically for children, not to mention household and public IoT products that also capture children's data, more research on children's privacy is essential, especially concerning:

- Emotion detection sentiment and analysis of children

- Child development in relation to IoT and AI toys

- A broadened exploration of children's privacy rights in light of new technology's ability to reveal more to parents and commercial actors

NORMS

Social norms are one of the most important factors in determining how privacy manifests in product development and how the public reacts to these products. Norms can either develop organically or be fostered. More research is needed on 'norm entrepreneurship' – the intentional shaping of norms by different actors. Research can help illuminate how:

- Governance actors such as State Attorneys General affect the privacy landscape[280]
- Commercial actors affect people's perception of privacy
- Journalism can affect people's privacy expectations
- Certification regimes, like the EU's suggestion of a 'Trusted IoT Label,' can affect privacy views and purchasing habits
- The increased presence of microphones and cameras in the home affect behavior

DISCRIMINATION

There is already a body of research emerging on the discriminatory effects of big data.[281] What's needed is additional research on how IoT devices are part of the supply chain of information that can lead to discriminatory practices, such as:

- Legal, commercial and technical research on data collected by wearables, by cars, or in the home leading to discriminatory insurance offerings
- Leakage of personal data from private contexts into the work environment leading to discriminatory firing or retribution, or effects on hiring practices
- The melding of IoT data with credit reports

280 Citron, 2016

281 See, e.g., Gangadharan, et al., 2014, "Data and Discrimination: Collected Essays": https://newamerica.org/documents/945/data-and-discrimination.pdf

INSURANCE

The cyber insurance market is young and usually focused on security issues, such as ransom and data breaches, but insurance policies for privacy-related failures are beginning to emerge. More research is needed on the intersection of privacy, insurance, governance, norms, and firm behavior. Research into privacy insurance for harms beyond identity theft for individuals, such as reputational damage from a data breach, would be valuable.

GOVERNANCE

Research into governance activities will help to show which methods work and which need improvement. Critical areas include:

- Further comparative analysis and interdisciplinary research on the lessons learned from food and drug safety, environmental harm, and product liability as models for governmental and private sector approaches to IoT privacy
- Further research into which private or government departments are best positioned to make privacy risk management an element of merger reviews
- Further research into the privacy risk calculus performed by shareholders during mergers and acquisitions
- Research into how Cost-Benefit Analysis can align with long-term social harms, such as the diminishment of private spaces or chilling effects on behavior and speech
- Continued legal research on IoT privacy in the context of product liability; this is especially important in light of the IoT's abilities to affect the physical environment
- Research into the effects of the GDPR's class action rights

FUNDING

We recommend allocating more funding to the following areas in order to enhance grantmakers' portfolios of privacy interventions and programs:

- Programs to embed more technologists with policymakers, such as the Congressional Innovation Fellowship program and the Ford-Mozilla Open Web Fellows program

- Privacy advocates and consumer protection groups, with an emphasis on technology issues for the latter

- Journalists, both to ensure better training on new technologies, and to sustain adversarial journalistic organizations like ProPublica

- Bridging identity management into other fields – IDM is barely represented in the academic community

- Projects that help consumers better understand the privacy and security aspects of IoT products, such as Consumer Reports' Digital Standard

- Creating visualizations and graphic communications of:
  > data flows for non-specialist audiences
  > representations of privacy harms and values

FOSTERING DIALOGUE

Governments, nonprofits and advocacy organizations have a continuing role to play in fostering dialogue, both in professional communities and in society at large. Dedicated educational outreach, public service announcements, and discussions at relevant conferences should be conducted on the following topics:

- The state of public and private surveillance
- The discourse of privacy values rather than harms: healthy democracy, full development of one's personality, decision-making free from manipulation

In addition, we encourage:

- More dialogue between the risk management community and the academic privacy community
- More dialogue between the insurance community and the privacy community
- Active public campaigns to enhance people's views of the value of their data
- More global dialogue to learn from different privacy cultures and governance models
- Professional dialogue to harmonize international risk frameworks and best practices
- Professional dialogue to further develop metrics, models, data sets, and testbeds to test and evaluate IoT products and best practices
- Publishing of case studies of IoT product privacy development

GOVERNANCE

The following are recommendations for IoT governance, broadly construed:

- Apply pressure on sector-specific legislation to include strong privacy and fairness elements
- Help create enhanced methods for American users to report privacy infringements; similar to Europeans' ability to report companies to data protection authorities
- Champion policies that require IoT manufacturers to indicate how long they will provide security updates for products
- Mandate data disclosure, transparency, and reporting to understand residual privacy risk from the IoT and manage this risk appropriately; e.g., expand laws to cover entities not covered by existing breach disclosure laws, require disclosure of breach events irrespective of whether a harm has been manifest, require more detailed and standardized breach reports, develop repositories of privacy breach incidents to help companies with situational awareness, require breach reporting within 72 hours
- Facilitate cyber insurance markets
- Actively enhance privacy norms at the government level through progressive procurement policies
- Experiment with tax, subsidy, and permit-based mechanisms to incentivize companies to make more socially desirable choices about privacy 'pollution'

RESEARCH PARTICIPANTS

The following people graciously lent their time and voices to this research. Their titles are the ones they had at the time of their participation.

Interviews

Justin Brookman, Director, Privacy and Technology Policy, Consumers Union

Michelle De Mooy, Director, Privacy & Data Project, Center for Democracy & Technology

Ben Dean, Ford/Media Democracy Fund Technology Exchange Fellow, Center for Democracy & Technology

Ian Glazer, VP Identity Product Management, Salesforce

Noah Harlan, Board Member, EdgeX Foundry

Joe Jerome, Policy Counsel, Privacy & Data Project, Center for Democracy & Technology

Meg Leta Jones, Assistant Professor, Communication, Culture & Technology, Georgetown University

Eve Maler, VP Innovation & Emerging Technology, ForgeRock

Andrew McStay, Professor of Digital Life, Bangor University

Dawn Nafus, Anthropologist, Intel

Heather Patterson, Senior Research Scientist, Privacy & Ethics, Intel

Rafi Rich, CEO, S.U.iT.S. (Smarter Urban iT & Strategies)

Florian Schaub, Assistant Professor, School of Information, University of Michigan

Lee Tien, Senior Staff Attorney, Electronic Frontier Foundation

Joris van Hoboken, Professor of Law, Vrije Universiteit Brussel

Representative #1, IoT Industry Company

Representative #2, IoT Industry Company

Senior US Government Official

Workshop Participants

Jim Adler, VP, Toyota Research Institute

Stephen Balkam, CEO, Family Online Safety Institute

Justin Brookman, Director, Privacy and Technology Policy, Consumers Union

Dan Caprio, Chairman, Providence Group

Betsy Cooper, Executive Director, Center for Long-Term Cyber Security, UC Berkeley

Ben Dean, Ford/Media Democracy Fund Technology Exchange Fellow, Center for Democracy & Technology

Laura Denardis, Professor, School of Communication, American University

Justin Erlich, former Principal Advisor on technology, data, and privacy to California Attorney General, California Department of Justice

Andrea Glorioso, Counselor (Digital Economy/Cyber), Delegation of the European Union to the United States

Stacey Gray, Policy Counsel, Future of Privacy Forum

Joe Jerome, Policy Counsel, Privacy & Data Project, Center for Democracy & Technology

Jennifer King, Co-Director, Center for Technology, Society & Policy, UC Berkeley

Emily McReynolds, Program Director, Tech Policy Lab, University of Washington

Melanie Millar-Chapman, Manager, Research, Office of the Privacy Commissioner of Canada

Maria Rerecich, Director, Electronics Testing Team, Consumers Union

Jatinder Singh, Senior Research Associate, University of Cambridge

Adam Thierer, Senior Research Fellow, Mercatus Center, George Mason University

Lee Tien, Senior Staff Attorney, Electronic Frontier Foundation

Ian Wallace, Senior Fellow in the International Security Program & Co-Director of the Cybersecurity Initiative, New America

Danny Weiss, Regional Director, Washington DC Area, Common Sense Media

Timothy Yim, Director of Data & Privacy, Startup Policy Lab

Meg Young, Ph.D Candidate, Information School and Tech Policy Lab, University of Washington

Advocate Technologist

Representative, IoT Industry Company

Representative, Social Media Company

Senior US Government Official

Senior Security Researcher

BIBLIOGRAPHY

Aberbach, J. and Rockman, B. (2002). Conducting and Coding Elite Interviews. *PS: Political Science and Politics, Vol. 35*(4). 673-676.

Ackerman, L. (2013). Mobile Health and Fitness Applications and Information Privacy. Privacy Rights Clearinghouse. Retrieved from https://www.privacyrights.org/sites/default/files/mobile-medical-apps-privacy-consumer-report.pdf

Agre, P. and Rotenberg, M. (1997). *Technology and Privacy: The New Landscape*. Cambridge: MIT Press.

Akalu, R., et al. (2016). Paving the way for Intelligent Transport Systems (ITS): The Privacy Implications of Vehicle Infotainment Systems. *Proceedings of the 6th ACM Symposium on Development and Analysis of Intelligent Vehicular Networks and Applications* (pp. 25-31). New York: ACM.

Altman, I. (1976). Privacy: A Conceptual Analysis. *Environment and Behavior, 8*(1): 7-30.

Altman, I. (1977). Privacy Regulation: Culturally Universal or Culturally Specific? *Journal of Social Issues 33*(3): 66-84.

Altman Vilandrie & Company. (2017). Are your company's IoT devices secure? Internet of Things Breaches are Common, Costly for U.S Firms. Retrieved from http://www.altvil.com/wp-content/uploads/2017/06/AVCo.-IoT-Security-White-Paper-June-2017.pdf

Antonakakis, et al. (2017). Understanding the Mirai Botnet. Proceedings of the 26th USENIX Security Symposium. Retrieved from https://www.usenix.org/system/files/conference/usenixsecurity17/sec17-antonakakis.pdf

Article 29 Data Protection Working Party. (2014). Opinion 8/2014 on the on Recent Developments on the Internet of Things. Retrieved from http://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2014/wp223_en.pdf

Article 29 Data Protection Working Party. (2017). Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679. Retrieved from http://ec.europa.eu/newsroom/document.cfm?doc_id=44137

Article 29 Data Protection Working Party. (2018). Guidelines on transparency under Regulation 2016/679. Retrieved from http://ec.europa.eu/newsroom/article29/document.cfm?action=display&doc_id=51025

Barth A., Datta, A., Mitchell, J., and Nissenbaum, H. (2006). Privacy and Contextual Integrity : Framework and Applications. *Proceedings of the IEEE Symposium on Security and Privacy.* Retrieved from https://ssrn.com/abstract=2567438

Bennett, C. and Raab, C. (2003). The Governance of Privacy: Policy Instruments in Global Perspective. Burlington: Ashgate Publishing.

Braun, V. and Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology, 3*(2), 77-101.

Brookman, J. and Hans, G.S. (2013). "Why Collection Matters: Surveillance as a De Facto Harm," Center for Democracy and Technology, available at http://www.futureofprivacy.org/wp-content/uploads/Brookman-Why-Collection-Matters.pdf

Bruening, P. and Patterson, H. (2016). A Context-Driven Rethink of the Fair Information Practice Principles. Retrieved from https://ssrn.com/abstract=2843315

California Department of Justice. (2013). Privacy on the Go: Recommendations for the Mobile Ecosystem. Retrieved from https://oag.ca.gov/sites/all/files/agweb/pdfs/privacy/privacy_on_the_go.pdf

Calo, R. (2012). Against Notice Skepticism in Privacy (And Elsewhere). *Notre Dame Law Review, 87*(3): 1027-1072.

Calo, R. (2013). Digital Market Manipulation. *George Washington Law Review, 82*(4): 995-1051.

Camp, J., Henry, R., Meyers, S., and Russo, G. (2016). Comment in response to Dept. of Commerce NTIA "Request for Comments on the Benefits, Challenges, and Potential Roles for the Government in Fostering the Advancement of the Internet of Things". Retrieved from https://www.ntia.doc.gov/files/ntia/publications/camp_et_al.pdf

Camp, L. and Johnson, M. (2012). *The Economics of Financial and Medical Identity Theft.* New York: Springer.

Campaign for a Commercial-Free Childhood. (n.d.). Hell No Barbie: 8 reasons to leave Hello Barbie on the shelf. Retrieved from http://www.commercialfreechildhood.org/action/hell-no-barbie-8-reasons-leave-hello-barbie-shelf

Castro D. and New, J. (2016). Everything the U.S. Government Is Doing to Help the Private Sector Build the Internet of Things. Center for Data Innovation. Retrieved from http://www2.datainnovation.org/2016-federal-support-iot.pdf

Cate, F., Cullen, P. and Mayer-Schönberger, V. (2014). Data Protection Principles for the 21st Century: Revising the 1980 OECD Guidelines. Retrieved from https://www.oii.ox.ac.uk/archive/downloads/publications/Data_Protection_Principles_for_the_21st_Century.pdf

Center for Information Policy Leadership. (2017). Comments by the Centre for Information Policy Leadership on the Article 29 Data Protection Working Party's "Guidelines on Transparency". Retrieved from https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_response_to_wp29_guidelines_on_transparency-c.pdf

Chang, A. (2018). The Facebook and Cambridge Analytica scandal, explained with a simple diagram. *Vox.* Retrieved from https://www.vox.com/policy-and-politics/2018/3/23/17151916/facebook-cambridge-analytica-trump-diagram

Citron, D. (2016). The Privacy Policymaking of State Attorneys General. *Notre Dame Law Review 92*(2): 747-816.

Cloud Security Alliance Mobile Working Group. (2015). Security Guidance for Early Adopters of the Internet of Things. Retrieved from https://downloads.cloudsecurityalliance.org/whitepapers/Security_Guidance_for_Early_Adopters_of_the_Internet_of_Things.pdf

Cogeco. (2017). The Cost of DDoS Attacks and Building the Business Case for Protection. Retrieved from https://www.cogecopeer1.com/wp-content/uploads/2017/03/Counting-the-Costs-of-DDoS-Attacks-DDoS-Services-Whitepaper.pdf

Cohen, J. (2013). What Privacy is For. *Harvard Law Review, 126*(7): 1904-1933.

Cohen, J. (2014). *The Private Life: Why We Remain in the Dark*. London: Granta Books.

Consumer Watchdog. (2017). Google, Amazon Patent Filings Reveal Digital Home Assistant Privacy Problems. Retrieved from http://www.consumerwatchdog.org/report/home-invasion-google-amazon-patent-filings-reveal-digital-home-assistant-privacy-problems

Costa, L. (2016). Privacy and the Precautionary Principle. *Computer Law & Security Review 28*(1): 14-24.

Dean, B. (2018). Strict Product Liability and the Internet of Things. Center for Democracy & Technology. Retrieved from https://cdt.org/files/2018/04/2018-04-16-IoT-Strict-Products-Liability-FNL.pdf

DeCew, J. (2013). Privacy. In E. Zalta (ed.). *Stanford Encyclopedia of Philosophy.* Stanford: The Metaphysics Research Lab.

Draper, K. (2018, Mar 13). Madison Square Garden Has Used Face-Scanning Technology on Customers. *New York Times.* Retrieved from https://www.nytimes.com/2018/03/13/sports/facial-recognition-madison-square-garden.html

Edara, K. (2014). *Keyword Determinations from Voice Data*. US2014/0337131A1

Epps, D., Saunders, K., and Tye, M. (2016). Verizon's Comments. Retrieved from https://www.ntia.doc.gov/files/ntia/publications/vz_comments_re_ntia_iot_notice_6-2.pdf

European Commission. (2016a). Advancing the Internet of Things in Europe. Retrieved from http://eur-lex.europa.eu/legal-content/TXT/?uri=CELEX:52016SC0110

European Commission. (2016b). Digital Single Market – Digitising European Industry Questions and Answers. Retrieved from http://europa.eu/rapid/press-release_MEMO-16-1409_en.htm

European Commission. (2017a). Cybersecurity Act. Retrieved from https://ec.europa.eu/transparency/regdoc/rep/1/2017/EN/COM-2017-477-F1-EN-MAIN-PART-1.PDF

European Commission. (2017b). Regulation on Privacy and Electronic Communications. Retrieved from http://ec.europa.eu/newsroom/dae/document.cfm?doc_id=41241

European Union. (2012). Charter of Fundamental Rights of the European Union 2012/C 326/02

Eversole, A., Day, T. and Brown, M. (2016). Comments of the US Chamber of Commerce Center for Advanced Technology and Innovation. Retrieved from https://www.ntia.doc.gov/files/ntia/publications/cati.iotcommentsfinal.pdf

Fadell, A., Matsuoka, Y., Sloo, D. and Veron, M. (2016). *Smart Home Automation System that Suggests or Automatically Implements Selected Household Policies Based on Sensed Observations*. US2016/0259308A1

Fairfield, J. and Engel, C. (2015). Privacy as a Public Good. *Duke Law Journal, 65*(3), 385-457.

Federal Trade Commission. (1980). FTC Policy Statement on Unfairness. Retrieved from https://www.ftc.gov/public-statements/1980/12/ftc-policy-statement-unfairness

Federal Trade Commission. (2015). Internet of Things: Privacy & Security in a Connected World. Retrieeved from https://www.ftc.gov/system/files/documents/reports/federal-trade-commission-staff-report-november-2013-workshop-entitled-internet-things-privacy/150127iotrpt.pdf

Federal Trade Commission. (2016a). Putting Disclosures to the Test: Staff Summary. Retrieved from https://www.ftc.gov/system/files/documents/reports/putting-disclosures-test/disclosures-workshop-staff-summary-update.pdf

Federal Trade Commission. (2017). Informational Injury Workshop. [slides] Retrieved from https://www.ftc.gov/system/files/documents/public_events/1256463/informational_injury_workshop_slides.pdf

Finch, K. (2016). A Visual Guide to Practical Data De-Identification. Future of Privacy Forum. Retrieved from https://fpf.org/2016/04/25/a-visual-guide-to-practical-data-de-identification/
Froomkin, A.M. (2015). Regulating Mass Surveillance as Privacy Pollution: Learning from Environmental Impact Statements. *University of Illinois Law Review 2015*(5): 1713-1790.

Foulon, J., Lanoie, P., and Laplante, B. (2002). Incentives for Pollution Control: Regulation or Information?, *Journal of Environmental Economics and Management, 44*(1): 169-187.

Garfinkel, S. (2015). NISTIR 8053: De-Identification of Personal Information. Retrieved from https://nvlpubs.nist.gov/nistpubs/ir/2015/NIST.IR.8053.pdf

Goodman, E. (2015). *The Atomic Age of Data Policies for the Internet of Things*. Queenstown: Aspen Institute.

Gray, S. (2016). Always On: Privacy Implications of Microphone-Enabled Devices. Future of Privacy Forum. Retrieved from https://fpf.org/wp-content/uploads/2016/04/FPF_Always_On_WP.pdf

Haggerty K., and Ericson, R. (2000). *British Journal of Sociology, 51*(4): 605–622.

Hirsch, D., (2006). Protecting the Inner Environment: What Privacy Regulation Can Learn from Environmental Law. *Georgia Law Review, 4*(1): 1-63.

HITRUST. (n.d.). HITRUST De-Identification Framework. Retrieved from https://hitrustalliance.net/de-identification/

IoT Security Foundation (n.d.) Best Practice User Mark FAQ and Terms of Use. Retrieved from https://www.iotsecurityfoundation.org/best-practice-user-mark/

Jones, M. (2014). Privacy Without Screens & the Internet of Other People's Things. *Idaho Law Review, 51*(3): 639-660.

Jones, M. and Meurer, K. (2016). Can (and Should) Hello Barbie Keep a Secret? *2016 IEEE International Symposium on Ethics in Engineering, Science and Technology.* Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2768507

Kaminsky, M. (2015). Robots in the Home: What Will We Have Agreed To? *Idaho Law Review 51*(3): 661-667.

Karpf, D. (2017). Sorting through the Facebook-Russia-Trump Advertising Story. Retrieved from https://medium.com/@davekarpf/sorting-through-the-facebook-russia-trump-advertising-story-d096e3df3edb

Kearney J., and Reynolds, A. (2016). Comments of the Consumer Technology Association. Retrieved from: https://www.ntia.doc.gov/files/ntia/publications/cta_comments_re_ntia_iot_rfc-final-060216_2.pdf

Kohnstamm, J. and Madhub, D. (2014). Mauritius Declaration on the Internet of Things. 36th *International Conference of Data Protection and Privacy Commissioners.* Retrieved from https://edps.europa.eu/sites/edp/files/publication/14-10-14_mauritius_declaration_en.pdf

Könings, B. and Schaub, F. (2011). Territorial Privacy in Ubiquitous Computing. *Eighth International Conference on Wireless On-Demand Network Systems and Services*. (pp. 104-108). New York: IEEE

Könings, B., Schaub, F., and Weber, M. (2016). Privacy and Trust in Ambient Intelligent Environments. In S. Ultes, F. Nothdurft, T. Heinroth, and W. Minker (Eds.). Next Generation Intelligent Environments. Dordrecht: Springer.

Koops, B-J, Newell, B., Timan, T., Škorvánek, I., Chokrevski, T., and GaliĐ, M. (2017). A Typology of Privacy. *University of Pennsylvania Journal of International Law 38*(2): 483-575.

Kyllo v. U.S., 533 U.S. 27 (2001).

Lindblom, C. (1979). Still Muddling: Not Yet Through. *Public Administration Review 39*(6): 517-526.

Lyon,D. (2008). Surveillance Society. Talk for Festival del Diritto, Piacenza, Italy. Retrieved from http://www.festivaldeldiritto.it/2008/pdf/interventi/david_lyon.pdf

Maheshwari, A. (2018, March 31). Hey, Alexa, What Can You Hear? And What Will You Do With It? *New York Times.* Retrieved from https://www.nytimes.com/2018/03/31/business/media/amazon-google-privacy-digital-assistants.html

Maldoff, G. (2016). Top 10 Operational Impacts of the GDPR: Part 6 – RTBF and data portability. IAPP: The Privacy Advisor. [blog]. Retrieved from https://iapp.org/news/a/top-10-operational-impacts-of-the-gdpr-part-6-rtbf-and-data-portability/

Maler, E. (2017). Designing a New Consent Strategy for Digital Transformation. RSA Conference. Retrieved from https://www.rsaconference.com/events/us17/agenda/sessions/6826-designing-a-new-consent-strategy-for-digital

Mallet, S. (2004). Understanding home: a critical review of the literature. *The Sociological Review 52*(1), 62-89.

Martin, E.D. (1928). Education. In C. Beard (ed.), *Whither Mankind: A Panorama of Modern Civilization*. New York: Longman, Green and Co.
Marx, G. (2012). Privacy Is Not Quite Like the Weather. In D. Wright and P. De Hert (Eds.). *Privacy Impact Assessment.* Dordrecht: Springer.

Massey, C., Sharkey, S., Roland, J., and Crawford, D. (2016). Comments of T-Mobile USA, Inc. Retrieved from https://www.ntia.doc.gov/files/ntia/publications/t-mobile_ntia_internet_of_things_comments_vfinal.pdf

McDonald, A. and Cranor, L. (2008). The Cost of Reading Privacy Policies. *I/S: A Journal of Law and Policy for the Information Society 4*(3): 543-565.

McKay, K., Bassham, L., Turan, M., and Mouha N. (2017). NISTIR 8114: Report on Lightweight Cryptography. Retrieved from https://nvlpubs.nist.gov/nistpubs/ir/2017/NIST.IR.8114.pdf

McStay, A. (2016). Empathic media and advertising: Industry, policy, legal and citizen perspectives (the case for intimacy). *Big Data & Society 3*(2): 1-11.

McStay, A. (2018). *Emotional AI: The Rise of Empathic Media.* London: Sage.

Montgomery, K., Chester, J., and Kopp, K. (2017). Health Wearable Devices in the Big Data Era: Ensuring Privacy, Security, and Consumer Protection. Center for Digital Democracy. Retrieved from https://www.democraticmedia.org/sites/default/files/field/public/2016/aucdd_wearablesreport_final121516.pdf

National Highway Traffic Safety Administration. (n.d.) Automated Driving Systems: FAQ: Voluntary Guidance. Retrieved from https://www.nhtsa.gov/manufacturers/automated-driving-systems#automated-driving-systems-faq

National Highway Traffic Safety Administration. (2016). Federal Automated Vehicles Policy. Retrieved from https://www.transportation.gov/sites/dot.gov/files/docs/AV%20policy%20guidance%20PDF.pdf

Nissenbaum, H. (2009) *Privacy in Context*. Stanford: Stanford University Press.

North, D. (1991). Institutions. *The Journal of Economic Perspectives, 5*(1): 97-112.

Office of the Privacy Commissioner of Canada. (2016). The Internet of Things: An introduction to privacy issues with a focus on the retail and home environments. Retrieved from https://www.priv.gc.ca/en/opc-actions-and-decisions/research/explore-privacy-research/2016/iot_201602/

Ohm, P. (2010). Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization. *UCLA Law Review, Vol. 57*(6): 1701-1777.

Ovum. (2015). Understanding the IoT Opportunity: An Industry Perspective – Use cases and opportunities in different verticals. Retrieved from https://www.linkedin.com/pulse/ovum-report-delivers-eight-clear-messages-internet-things-niall-hunt/

Palen, L. and Dourish, P. (2003). Unpacking "privacy" for a networked world. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 129-136). New York: ACM.

Patterson, H. (2013). Contextual Expectations of Privacy in Self-Generated Health Information Flows. *TPRC 41: The 41st Research Conference on Communication, Information and Internet Policy.* Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2242144

Peppet, S. (2014). Regulating the Internet of Things: First Steps Toward Managing Discrimination, Privacy, Security, and Consent. *Texas Law Review 93*(1): 85-176.

Pew Research Center. (2014). The Internet of Things Will Thrive By 2025. Retrieved from http://www.pewinternet.org/2014/05/14/internet-of-things/

Picard, R. (1995). Affective Computing. M.I.T Media Laboratory Perceptual Computing Section Technical Report No. 321. Retrieved from https://affect.media.mit.edu/pdfs/95.picard.pdf

Powles, J. (2015, Mar 11). We are citizens, not mere physical masses of data for harvesting. *The Guardian*. Retrieved from https://www.theguardian.com/technology/2015/mar/11/we-are-citizens-not-mere-physical-masses-of-data-for-harvesting

President's Council of Advisors on Science and Technology. (2014). Big Data and Privacy: A Technical Perspective. Retrieved from https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/PCAST/pcast_big_data_and_privacy_-_may_2014.pdf

Prins, C. (2006). When personal data, behavior and virtual identities become a commodity: Would a property rights approach matter? *SCRIPT-ed, 3*(4): 270-303.

Prosser, W. (1960). Privacy. *California Law Review 48*(3): 383-423.

Regan, P. (1995). *Legislating Privacy. Technology, Social Values, and Public Policy.* Chapel Hill: University of North Carolina Press.

Reidenberg, J. (2014). Privacy in Public. *University of Miami Law Review 69*(1): 141-160.

Rosner, G. (2016a). Internet of Things Privacy Forum Comments Submitted to the National Telecommunications and Information Administration. Retrieved from https://www.ntia.doc.gov/files/ntia/publications/rosner_response_-_iot_privacy_forum.pdf

Rosner, G. (2016b). In the age of connected devices, will our privacy regulations be good enough? *O'Reilly*. [blog]. Retrieved from https://www.oreilly.com/ideas/in-the-age-of-connected-devices-will-our-privacy-regulations-be-good-enough

Sadeh, N. (2017). Privacy in the Age of IoT: Technologies to Help Users, Developers, IoT Vendors and Regulators. Presentation at Privacy Engineering Research and the GDPR: A Trans-Atlantic Initiative. Retrieved from https://www.cylab.cmu.edu/_files/pdfs/partners/conference2017/Sadeh.pdf

Saldaña, J. (2009). *The Coding Manual for Qualitative Researchers*. London: Sage.

Santucci, G. (2011). The Internet of Things: The Way ahead. In O. Vermesan and P. Friess (eds.) *Internet of Things - Global Technological and Societal Trends From Smart Environments and Spaces to Green ICT*. Aalborg: River Publishers.

Santucci, G. (2013). Privacy in the Digital Economy: Requiem or Renaissance? An Essay on the The Future of Privacy. The Privacy Surgeon. Retrieved from http://www.privacysurgeon.org/blog/wp-content/uploads/2013/09/Privacy-in-the-Digital-Economy-final.pdf

Sarnacki, D. (1984). Analyzing the Reasonableness of Bodily Intrusions. *Marquette Law Review, 68*(1), 130-153.

Schaub, F., Balabako, R., Durity, A., and Cranor L. (2015). A Design Space for Effective Privacy Notices. *Proceedings of the Eleventh Symposium on Usable Privacy and Security (SOUPS 2015)* Retrieved from https://www.usenix.org/system/files/conference/soups2015/soups15-paper-schaub.pdf

Schifferle, L. (2016). What is your phone telling your rental car? FTC Consumer Information. [blog]. Retrieved from https://www.consumer.ftc.gov/blog/2016/08/what-your-phone-telling-your-rental-car

Schneier, B. (2016, Jan 26). IoT Security: The Market Has Failed. Schneier on Security. Retrieved from https://www.schneier.com/news/archives/2017/01/iot_security_the_mar.html

Schwartz, P. (1999). Internet Privacy and the State. *Connecticut Law Review, 32*: 815-859.

Schwartz, P. (2004). Property, Privacy, and Personal data. *Harvard Law Review, 117*(7): 2056-2128.

Schwartz, P. (2013). Information Privacy in the Cloud. *University of Pennsylvania Law Review 161*(6): 1623-1662.

Scott, W. (2003). Institutional carriers: reviewing modes of transporting ideas over time and space and considering their consequences. *Industrial and Corporate Change, 12*(4): 879-894.

Sholtz, P. (2001). Transactions Costs and the Social Costs of Online Privacy. *First Monday, 6*(5). Retrieved from http://www.ojphi.org/ojs/index.php/fm/article/view/859/768

Sikdar, A., Behera, S., and Dogra, D. (2016). Computer-Vision-Guided Human Pulse Rate Estimation: A Review. *IEEE Reviews in Biomedical Engineering, 9*, 91-105.

Slomovic, A. (2015). Workplace Wellness, Privacy and the Internet of Things. [blog post]. Retrieved from https://www.annaslomovic.com/single-post/2015/01/31/Workplace-Wellness-Privacy-and-the-Internet-of-Things

Singer, N. (2016 Feb 28). Why a Push for Online Privacy Is Bogged Down in Washington. *New York Times*. Retrieved from https://www.nytimes.com/2016/02/29/technology/obamas-effort-on-consumer-privacy-falls-short-critics-say.html

Sloan, R. and Warner, R. (2013). Beyond Notice and Choice: Privacy, Norms, and Consent. *Journal of High Technology Law 14*(2): 370-412.

Solove, D. (2006). A Taxonomy of Privacy. *University of Pennsylvania Law Review, 154*(3), 477-560.

Stewart, R. (2002). Environmental regulatory decision making under uncertainty. In T. Swanson (Ed.). *An Introduction to the Law and Economics of Environmental Policy: Issues in Institutional Design.* Bingley: Emerald Publishing

Sunstein, C. (2003). Beyond the Precautionary Principle. *University of Pennsylvania Law Review, 151*(3): 1003-1058.

Tene, O. (2011). Privacy: The New Generations. *International Data Privacy Law, 1*(1): 15–27.

The Economist. (2018, Mar 22). A pedestrian has been killed by a self-driving car. *The Economist.* Retrieved from https://www.economist.com/news/science-and-technology/21739149-driverless-tragedy-pedestrian-has-been-killed-self-driving-car

Thierer, A. (2013). Technopanics, Threat Inflation, and the Danger of an Information Technology Precautionary Principle. *Minnesota Journal of Law, Science & Technology, 14*(1): 309-386.

Thierer, A. (2016). *Permissionless Innovation: The Continuing Case for Comprehensive Technological Freedom.* Arlington: Mercatus Center at George Mason University.

UK Information Commissioner's Office. (n.d.) Where should you deliver privacy information to individuals? Retrieved from https://ico.org.uk/for-organisations/guide-to-data-protection/privacy-notices-transparency-and-control/where-should-you-deliver-privacy-information-to-individuals/

UK Information Commissioner's Office. (2016). Privacy regulators study finds Internet of Things shortfalls. Retrieved from https://ico.org.uk/about-the-ico/news-and-events/news-and-blogs/2016/09/privacy-regulators-study-finds-internet-of-things-shortfalls/

UN General Assembly. (1992). Rio Declaration on Environment and Development. Retrieved from http://www.un.org/documents/ga/conf151/aconf15126-1annex1.htm

US Department of Commerce. (2017). Fostering the Advancement of the Internet of Things. Retrieved from https://www.ntia.doc.gov/files/ntia/publications/iot_green_paper_01122017.pdf

US Department of Health and Human Services. (2016). Examining Oversight of the Privacy & Security of Health Data Collected by Entities Not Regulated by HIPAA. Retrieved from https://www.healthit.gov/sites/default/files/non-covered_entities_report_june_17_2016.pdf

US Department of Health, Education and Welfare. (1973). Records, Computers and the Rights of Citizens. Retrieved from https://epic.org/privacy/hew1973report/default.html

US Department of Homeland Security. (2008). The Fair Information Practice Principles. Memorandum Number 2008-01. Retrieved from https://www.dhs.gov/sites/default/files/publications/privacy_policyguide_2008-01_0.pdf

US Environmental Protection Agency. (n.d.). Economic Incentives. Retrieved from https://www.epa.gov/environmental-economics/economic-incentives.

Uteck, A. (2013). Reconceptualizing Spatial Privacy for the Internet of Everything. (Unpublished doctoral dissertation). University of Ottawa, Ottawa, Ontario.

Vaismoradi, M., Jones, J., Turunen, H., and Snelgrove, S. (2016). Theme development in qualitative content analysis and thematic analysis. *Journal of Nursing Education and Practice, 6*(5). 100-110.

Van den Hoven, J., Guimarães Pereira, Â., Dechesne, F., Timmermans, J., and Vom Lehn, H. (2016). Fact sheet – Ethics Subgroup IoT – version 4.0. Retrieved from http://ec.europa.eu/information_society/newsroom/cf/dae/document.cfm?doc_id=1751

Vila, T., Greenstadt, R., and Molnar, D. (2003). Why We Can't be Bothered to Read Privacy Policies: Models of Privacy Economics as Lemons Market. *Proceedings of the 5th International Conference on Electronic Commerce.* Retrieved from http://www.dmolnar.com/papers/econprivacy.pdf

Warren, S. and Brandeis, L. (1890). The Right to Privacy. *Harvard Law Review, 4*(5): 193-220.

Westin, A. 1967. *Privacy and Freedom.* New York: Atheneum.

Woolf, N. (2016 Oct 26). DDoS attack that disrupted internet was largest of its kind in history, experts say. *The Guardian.* Retrieved from https://www.theguardian.com/technology/2016/oct/26/ddos-attack-dyn-mirai-botnet

Zuboff, S. (2015). Big other: surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology 30*(1): 75–89.

APPENDIX 1
IOT PRIVACY-RELEVANT STANDARDS AND CONSORTIA

MULTI-LAYER CONSORTIA

IPSO Alliance	https://www.ipso-alliance.org/
Eclipse Foundation IoT	https://iot.eclipse.org/
oneM2M	http://www.onem2m.org/
EdgeX Foundry	https://www.edgexfoundry.org/
TeleHash	http://telehash.org/
IoTivity	https://www.iotivity.org/
OMA LightweightM2M	http://www.openmobilealliance.org/wp/overviews/lightweightm2m_overview.html
W3C Web of Things	https://www.w3.org/WoT/

IDENTITY MANAGEMENT

User-Managed Access Protocol	https://kantarainitiative.org/confluence/display/uma/Home
Consent Receipts	https://kantarainitiative.org/confluence/display/infosharing/Consent+Receipt+Specification
OAuth 2.0	https://oauth.net/2/
Health Relationship Trust (HEART)	https://openid.net/wg/heart/
Authorization and Authentication for Constrained Environments	https://tools.ietf.org/wg/ace/charters

STANDARDS DEVELOPMENT EFFORTS

ITU-T SG20 Internet of Things and Smart Cities and Communities	https://www.itu.int/en/ITU-T/studygroups/2013-2016/20/Pages/default.aspx
IEEE P1912 WG - Privacy and Security Architecture for Consumer Wireless Devices Working Group	https://standards.ieee.org/develop/wg/P1912_WG.html
ISO 29100 - Privacy Framework	https://www.iso.org/standard/45123.html

CONNECTED BODY

Personal Connected Health Alliance	http://www.pchalliance.org/
Wireless Life Sciences Alliance	http://wirelesslifesciences.org/

CONNECTED HOME

Thread Group	https://www.threadgroup.org/
Z-Wave	http://www.z-wave.com/
Insteon	https://www.insteon.com/
HomePlug	http://www.homeplug.org/
Home Gateway Initiative	http://www.homegatewayinitiative.org/
Weave	https://nest.com/weave/

CONNECTED BUILDINGS

Smart Buildings Alliance	http://www.smartbuildingsalliance.org/
EnOcean Alliance	https://www.enocean-alliance.org/
Lonworks	http://www.lonmark.org/

CONNECTED CARS AND TRANSPORTATION

Genivi	https://www.genivi.org/
Open Automotive Alliance	https://www.openautoalliance.net/#about
Automotive-grade Linux	https://www.automotivelinux.org/
IEEE 1609 - Wireless Access in Vehicular Environments	https://standards.ieee.org/develop/wg/1609.html
Cellular V2X	https://www.qualcomm.com/invention/technologies/lte/advanced-pro/cellular-v2x
SAE J2945/1 - On-Board System Requirements for V2V Safety Communications	https://www.sae.org/standards/content/j2945/1_201603/
IP Wireless Access in Vehicular Environments	https://tools.ietf.org/wg/ipwave/charters

APPENDIX 2
IOT PRIVACY-RELEVANT ACADEMIC CENTERS

SOCIAM: The Theory and Practice of Social Machines
http://sociam.org

The Information Society Project, Yale Law School
https://law.yale.edu/isp

Privacy Lab, Yale Law School
https://privacylab.yale.edu

Department of Science, Technology, Engineering and Public Policy,
University College London
http://www.ucl.ac.uk/steapp

Center on Privacy & Technology, Georgetown Law
http://www.law.georgetown.edu/academics/centers-institutes/privacy-technology/index.cfm

Center for Information Technology Policy, Princeton University
https://citp.princeton.edu/

Berkman Klein Center for Internet & Society, Harvard University
https://cyber.harvard.edu

Oxford Internet Institute, Oxford University
https://www.oii.ox.ac.uk

Digital Ethics Lab, Oxford University
http://digitalethicslab.oii.ox.ac.uk

Cyber Security Oxford, University of Oxford
https://www.cybersecurity.ox.ac.uk

Tech Policy Lab, University of Washington
http://techpolicylab.org

Connected & Open Research Ethics, UC San Diego
http://thecore.ucsd.edu

Center for Internet & Society, Stanford University
http://cyberlaw.stanford.edu

Pervasive Data Ethics for Computational Research, University of Maryland
https://pervade.umd.edu

Cyber Technology Institute, De Montfort University
http://www.dmu.ac.uk/research/research-faculties-and-institutes/technology/cyber-technology-institute/cti-home.aspx

Internet Ethics program at the Markkula Center for Applied Ethics,
Santa Clara University
https://www.scu.edu/ethics/focus-areas/internet-ethics/

Security and Privacy Research Group, University of Birmingham
http://sec.cs.bham.ac.uk

CyLab Security & Privacy Institute, Carnegie Mellon University
https://www.cylab.cmu.edu

Security and Privacy Research Lab, University of Washington
https://seclab.cs.washington.edu

Digital Policy Institute, Ball State University
http://www.digitalpolicyinstitute.org

Laboratory of Pervasive Computing, Tampere University
http://twitter.com/tampereunitech

Semaphore research cluster at the iSchool, University of Toronto,
http://semaphore.utoronto.ca

Pervasive Interaction Technology Lab, IT University of Copenhagen
http://pitlab.itu.dk

AI Now Institute, New York University
https://ainowinstitute.org

Data and Society, London School of Economics,
http://www.lse.ac.uk/study-at-lse/Graduate/Degree-programmes-2018/MSc-Media-and-Communications-Data-and-Society

Trustworthy Technologies Strategic Research Initiative, Cambridge University
https://www.trusttech.cam.ac.uk

Tilburg Institute for Law, Technology, and Society
http://www.tilburguniversity.edu/research/institutes-and-research-groups/tilt/

Centre for Information Rights, University of Winchester
https://www.winchester.ac.uk/research/building-a-sustainable-and-responsible-future/centre-for-information-rights/

Centre for Research into Information, Surveillance and Privacy (CRISP)
http://www.crisp-surveillance.com

The Nordic Centre for Internet and Society, Norwegian Business School
https://www.bi.edu/research/find-departments-and-research-centres/research-centres/nordic-centre-for-internet-and-society/

Meaningful Consent in the Digital Economy Project, University of Southampton
http://www.meaningfulconsent.org

Cyber-Physical Lab, Newcastle University,
https://research.ncl.ac.uk/cplab/

Resilient Cyber-Physical Systems Lab, University of Maryland
http://www.ece.umd.edu

Cyber Security Centre, University of Warwick
https://warwick.ac.uk/fac/sci/wmg/research/csc/

Cyber Security Research Center, Ben-Gurion University
https://cyber.bgu.ac.il

Interdisciplinary Research Centre in Cyber Security, University of Kent
https://cyber.kent.ac.uk

Centre for Cyber Security, University of Surrey
https://www.surrey.ac.uk/surrey-centre-cyber-security

Internet of Things and People Research Center, Malmö University
http://iotap.mah.se

Horizon Digital Economy Research Institute, University of Nottingham
https://www.horizon.ac.uk

Center for Long Term Cybersecurity, UC Berkeley
https://cltc.berkeley.edu

Data & Society Research Institute
https://datasociety.net

Sensor City, Liverpool John Moores University
https://www.ljmu.ac.uk/about-us/sensor-city

International Computer Science Institute, UC Berkeley
http://www.icsi.berkeley.edu/icsi/

New America Open Technology Institute
https://www.newamerica.org

APPENDIX 3
IOT PRIVACY-RELEVANT CONFERENCES

STRIVE 2018: First Intl. Workshop On Safety, Security, And Privacy
In Automotive Systems
Sep 18, 2018, Västerås, Sweden
http://www.iit.cnr.it/strive2018

PST 2018: The Sixteen International Conference on Privacy, Security and Trust
Aug 28-30, 2018, Belfast, UK
http://pstnet.ca/pst2018/

CENTRIC 2018: The Eleventh International Conference on Advances in Human-oriented and Personalized Mechanisms, Tech and Services
Oct 14-18, 2018, Nice, France
http://iaria.org/conferences2018/CENTRIC18.html

IEEE ISC2 2018: IEEE International Smart Cities Conference
Sep 16-19, 2018, Kansas City, MO
http://sites.ieee.org/isc2-2018/

TELERISE 2018: Fourth Intl. Workshop On Technical And Legal Aspects
Of Data Privacy And Security
Sep 2, 2018 Budapest, Hungary
http://www.iit.cnr.it/telerise2018/

CPDP2019: Computers, Privacy and Data Protection
Jan 30 – Feb 1, 2019, Brussels, Belgium
http://www.cpdpconferences.org

The 6th International Symposium on Security and Privacy on Internet of Things
Dec 12 - 15, 2017, Guangzhou, China
http://trust.gzhu.edu.cn/conference/SPIoT2017/

1st IEEE Intl Workshop on Intelligent Multimedia Applications and Design
for Quality Living
Dec 11 - 13, 2017, Taichung, Taiwan
http://ism2017.asia.edu.tw/imad-2017/

SITN 2018 The 2nd International Workshop on Securing IoT Networks
Jul 30 - Aug 3, 2018 Halifax, Canada
http://cse.stfx.ca/~iThings2018/download/SITN-2018%20CFP_ZT_v3.pdf

ICI 2018 The 5th International Symposium on Intercloud and IoT
Aug 6 - 8, 2018 Barcelona
http://www.ficloud.org/ici2018/

IoT-SECFOR 2018 2nd International Workshop on Security and Forensics of IoT
Aug 27 - 30, 2018, Hamburg, Germany
https://www.ares-conference.eu/workshops/iot-secfor-2018/

FTC PrivacyCon
F*eb 28, 2018, Washington DC*
https://www.ftc.gov/news-events/events-calendar/2018/02/privacycon-2018

Amsterdam Privacy Conference
Oct 5 - 8, 2018
http://apc2018.com

Cybersecurity 2017: The 6th International Workshop on Cyber Security and Privacy
Oct 12 - 14, 2017, Nanjing, China
http://cyberc.org/Program/Security

IEEE Symposium on Security and Privacy
May 20 - 24 May 2018, San Francisco, CA
http://www.ieee-security.org/TC/SP2018/

International Conference on Intelligent Information Technologies
Dec 20 - 22, 2017, Chennai, India
https://www.iciit.in/

ACM Symposium on Applied Computing: Special Track on Internet of Things
Apr 9 - 13, 2018 Pau, France
http://infohost.nmt.edu/~zheng/doku.php?id=sac2018

3rd International Conference on Internet of Things, Big Data and Security
M*ar 19 - 21, 2018, Funchal, Madeira, Portugal*
http://iotbds.org/

14th EAI International Conference on Security and Privacy in Communication Networks
August 8-10, 2018, Singapore
http://securecomm.org

IoT Tech Expo Global
April 18 - 19, 2018, London, UK
https://www.iottechexpo.com/global/

AmI 2018 European Conference on Ambient Intelligence
Nov 12 - 14, 2018, Golden Bay Beach Hotel, Larnaca, Cyprus
http://www.cyprusconferences.org/ami2018

The Eighth International Conference on Ambient Computing, Applications, Services and Technologies
Nov 18 - 22, 2018 Athens, Greece
http://iaria.org/conferences2018/AMBIENT18.html

ANT 2018 9th International Conference on Ambient Systems, Networks and Technologies
May 8 - 11, 2018 Porto, Portugal
http://cs-conferences.acadiau.ca/ant-18/

ICMU 2018 11th International Conference on Mobile Computing and Ubiquitous Networking
Oct 5 - 8, 2018, Auckland, New Zealand
http://www.icmu.org/icmu2018/

ISBA 2018 IEEE International Conference on Identity, Security and Behavior Analysis
Jan 10 - 12, 2018 Singapore
http://www3.ntu.edu.sg/conference/ISBA2018/home.htm

FiUSE 2018 International Workshop on Fog Computing in Internet of Things and Ubiquitous Systems Engineering
Jun 11 - 12, 2018, Tallinn, Estonia
https://fiuse2018.cs.ut.ee/

MMEDIA 2018 The Tenth International Conference on Advances in Multimedia
Apr 22 - 26, 2018 Athens, Greece
http://iaria.org/conferences2018/MMEDIA18.html

WONS 2018 14th Wireless On-demand Network systems and Services Conference
Feb 6 - 8, 2018 Isola, France
http://2018.wons-conference.org/

UBICOMM 2018 The Twelfth International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies
Nov 18 - Nov 22, 2018 Athens, Greece
http://iaria.org/conferences2018/UBICOMM18.html

MUE 2018 The 12th International Conference on Multimedia and Ubiquitous Engineering
Apr 23 -25, 2018, Salerno, Italy
http://www.mue-conference.org/2018

WristSense 2018: Workshop on Sensing Systems and Applications Using Wrist Worn Smart Devices
Mar 19, 2018, Athens, Greece
https://sites.google.com/view/wristsense2018

EUSPN 2018 9th International Conference on Emerging Ubiquitous Systems and Pervasive Networks
Nov 5 - 8, 2018, Leuven, Belgium
http://cs-conferences.acadiau.ca/euspn-18/

PICom 2018 16th IEEE International Conference on Pervasive Intelligence and Computing
Aug 12 - 15, 2018, Athens, Greece
http://cyber-science.org/2018/picom/

CoWPER 2018 Toward A City-Wide Pervasive EnviRonment
Jun 11, 2018, Hong Kong
http://secon2018.ieee-secon.org/the-cowper-workshop/

IUPT 2018 8th International Symposium on Internet of Ubiquitous and Pervasive Things
May 8 - 11, 2018, Porto, Portugal
https://hud-cs-research.github.io/iupt2018/

8th ACM MobiHoc 2018 Workshop on Pervasive Wireless Healthcare Workshop
Jun 25, 2018, Los Angeles, USA
https://www.sigmobile.org/mobihoc/2018/workshop-mobile-health.html

PerLS 2018 The Second International Workshop on Pervasive Smart Living Spaces
Mar 19 - 23, 2018, Athens, Greece
http://iotap.mah.se/perls2018/

PerIoT 2018 Second International Workshop on Mobile and Pervasive Internet of Things
Mar 19 - 23, 2018, Athens, Greece
https://periot.github.io/2018/

PerCom 2019 IEEE International Conference on Pervasive Computing and Communications
Mar 11-15, Kyoto, Japan
http://www.percom.org/

12th EAI International Conference on Pervasive Computing Technologies for Healthcare
May 21 - 24, 2018, New York, NY
http://pervasivehealth.org/

Layout & Design
Ian Estevens estevensian@gmail.com