



Information Systems Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Special Section Introduction—Ubiquitous IT and Digital Vulnerabilities

Sam Ransbotham, Robert G. Fichman, Ram Gopal, Alok Gupta

To cite this article:

Sam Ransbotham, Robert G. Fichman, Ram Gopal, Alok Gupta (2016) Special Section Introduction—Ubiquitous IT and Digital Vulnerabilities. Information Systems Research 27(4):834-847. <https://doi.org/10.1287/isre.2016.0683>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2016, INFORMS

Please scroll down for article—it is on subsequent pages

INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Special Section Introduction

Ubiquitous IT and Digital Vulnerabilities

Sam Ransbotham, Robert G. Fichman

Information Systems, Carroll School of Management, Boston College, Chestnut Hill, Massachusetts 02467
{sam.ransbotham@bc.edu, fichman@bc.edu}

Ram Gopal

Department of Operations and Information Management, University of Connecticut, Storrs, Connecticut 06269,
ram.gopal@business.uconn.edu

Alok Gupta

Carlson School of Management, University of Minnesota, Minneapolis, Minnesota 55455, alok@umn.edu

While information technology benefits society in numerous ways, it unfortunately also has potential to create new vulnerabilities. This special issue intends to stimulate thought and research into understanding and mitigating these vulnerabilities. We identify four mechanisms by which ubiquitous computing makes various entities (people, devices, organizations, societies, etc.) more vulnerable, including: increased visibility, enhanced cloaking, increased interconnectedness, and decreased costs. We use the papers in the special issue to explain these mechanisms, and then outline a research agenda for future work on digital vulnerabilities spanning four areas that are, or could become, significant societal problems with implications at multiple levels of analysis: Online harassment and incivility, technology-driven economic inequality, industrial Internet of Things, and algorithmic ethics and bias.

Keywords: online harassment; economic inequality; Internet of things; algorithmic bias; algorithmic ethics; artificial intelligence

History: Published online in *Articles in Advance* November 29, 2016.

1. Introduction

Information technology (IT) incites hyperbole. In a relatively short time, it has profoundly changed society—usually for the better but sometimes for the worse. IT now permeates practically every aspect of individual, organizational, social, and economic activity. A result of this ubiquity is to create new vulnerabilities that are not yet fully understood. While positive aspects of pervasive digitization attract considerable interest, this special issue focuses instead on vulnerabilities introduced or exacerbated by new (and old) forms of IT. How can we learn more about this darker side to better manage organizations, societies, and our personal lives?

IT adds value in numerous ways. Organizations can use systems to gather, organize, select, synthesize, and distribute information across all areas of a value chain including operations, logistics, production, marketing, design, service, and infrastructure. A plethora of emerging applications of IT provide significant opportunities for the creation of social and economic value.

Yet, there is no panacea; these same technologies can also increase our vulnerability. Emerging technologies frequently have unintended consequences, or may create one problem even as they solve another (Overby et al. 2010). We illustrate this tension by focusing

on four mechanisms—visibility, cloaking, interconnection, and cost—that tend to magnify vulnerabilities and their potential harms. For each mechanism, we describe potential positive and negative effects. Then we illustrate the resulting tension through papers in this special issue. These papers offer guidance for better managing organizations, societies, and our personal lives in light of these vulnerabilities. Finally, we highlight four important topics for future research.

2. Mechanisms Increasing the Prevalence of Vulnerabilities

We define digital vulnerability as a condition of susceptibility to harm that stems from the use of digital technologies. The harm can arise from the presence of the technology itself (e.g., onset of some condition such as technostress or digital addiction), an externality from another's use of the technology (e.g., inadvertent disclosure or data loss), or from intentional nefarious action (e.g., a security intrusion or privacy invasion). A vulnerability is *manifested* when the attack or condition actually occurs.¹

¹ For conciseness we will sometimes refer to digital vulnerabilities in terms of their manifestations, i.e., what some entity is *vulnerable to*.

What is it about the ubiquity of digital technology that enables or exacerbates so many vulnerabilities? We identify four general mechanisms by which ubiquitous computing makes various entities—e.g., people, organizations, societies, objects, systems, processes—more vulnerable. These are: (1) increased visibility, (2) enhanced cloaking, (3) increased interconnectedness, and (4) decreased costs. Interestingly, each of these mechanisms can also be linked to the positive side of ubiquity. This dialectic tension, in which the same mechanisms that drive value also magnify vulnerability, is the crux; it is what makes vulnerability mitigation such a challenge.

2.1. Increased Visibility

Any technology that enables the storage and retrieval of information about some entity—that is to say, just about any digital technology—enhances the potential visibility of that entity and its attributes. Even technologies intended just to automate a process or task simultaneously informate it, i.e., create a digital record of how that process unfolded (Zuboff 1994). This increased visibility has several facets: (a) breadth (a greater variety of attributes and qualities can now be seen), (b) depth (more detailed information is tracked), (c) reach (a larger, potentially unintended, audience of other entities can see the information), and (d) permanence (information is, in principle, available continuously and forever).

In the age of ubiquity we see a virtual explosion in the visibility of people, organizations, and physical objects. Through our use of mobile phones, wearables, Web browsers, social media, smart products, and many more technologies we create a continuous digital record of our traits, states, locations, behaviors, communications, and social graph. Digital surveillance and facial recognition technologies can reveal our identities to governments, businesses, or even private individuals without our knowledge. Internet search engines and social media make organizations and their products visible in ways they have never been before. Because of the increasing practice of embedding sensors, processors, and wireless networking into ever more physical objects, they too have become more visible.

The positive side of enhanced visibility is self-evident: organizations and products can serve us better if they can know more about us, and we can make better choices when we know more about them. “Smart” products can do lots of nice things that “dumb” ones cannot. Also, as a general rule we would like our organizations and institutions to be more transparent.

The negative side of enhanced visibility is that it can reveal information that people and organizations would prefer to keep—and often have a right to keep—hidden. Increased visibility enables theft of intellectual property (IP) and trade secrets, loss of privacy, identity theft, expanded legal discovery (Overby et al. 2010), and other manifestations of digital vulnerability.

In this special issue, Cavusoglu et al. (2016) play off the visibility tensions described above in their examination of Facebook. In December 2009, Facebook increased the granularity of privacy controls related to content sharing. For example, users could now restrict the visibility of wall posts just to certain audiences. While ostensibly a move to soothe the concerns of advocates about the privacy-related vulnerabilities created by Facebook, the potential of the change to *increase* sharing was noted by CEO Mark Zuckerberg, who said “When you have control over what you share, you want to share more” (Zuckerberg 2010). Panel data from over 13,000 active Facebook users supports this premise. In particular, after the change, the average Facebook user increases use of wall posts (which are more visible) and decreases use of private messages (which are less visible), although some interesting nuances arise that depend on an individual’s privacy sensitivity.

Lappas et al. (2016) examine visibility quite directly in their study of the reputational effects of injected fake reviews. These reviews enable rankings; rankings help individual find establishments that best match their needs and desires. When used as intended, better matches become more visible. However, firms are vulnerable to unwarranted shifts in their visibility due to the injection of fake reviews that either denigrate the firm or praise its competitors. The authors develop an operational measure of a firm’s visibility based on a firm’s relative position in the platform’s review-based ranking, and on the number of features desired by a user that a firm covers. Data from over 2.3 million reviews of 4,709 hotels shows that injection of only a small number of fake reviews (as few as 50) can significantly diminish a firm’s visibility. The authors also evaluate mitigation strategies for attacked hotels, such as detecting and disputing fake reviews.

2.2. Enhanced Cloaking

Even as digital technology makes things more visible, it also provides new ways for organizations and individuals to cloak their identities, characteristics, motivations, and behaviors. The fields of cybersecurity and privacy have been quite active in developing technologies to promote cloaking, such as encryption, privacy settings (Cavusoglu et al. 2016), and various other protection schemes, such as anonymous browsing. The positive side of these technologies is embodied in their explicit purpose: to enhance security and privacy by

So, we may refer to embarrassing exposures as a vulnerability, even though it would be more precise to say susceptibility to embarrassing exposures is a vulnerability, or alternatively, that embarrassing exposures are a manifestation of this vulnerability.

counteracting the effects of increased visibility. Also, and somewhat paradoxically, the ability to cloak some information about us, such as our identities, can promote greater sharing of other things about us, such as our ideas and opinions.

Bad actors have been equally vigorous in developing clever cloaking schemes to cover their tracks, such as IP spoofing, anonymous identities, and fake accounts. The anonymization underlying the Tor browser enables illegal commerce on the Dark Web just as easily as it empowers confidential communication. Thus there is also a negative side of cloaking, which is the enablement of security intrusion and cyber harassment, secret observation, digital impersonation, deceptions stemming from sock puppetry and fraudulent reviews (Lappas et al. 2016), and other manifestations of digital vulnerability.

Two articles from the special issue address vulnerabilities strongly linked to cloaking. Lowry et al. (2016) attend to the role of social anonymity as an enabler of cyberbullying on social media. The vulnerability of individuals to online harassment is a growing problem. In fact, a majority of Internet users report experiencing at least mild harassment (Duggan 2014). Lowry et al. (2016) develop a theory, based on an elaboration of Akers' (2011) social structure and social learning model, to explain why individuals engage in cyberbullying on social media. Based on data gathered from 1,002 adult Internet users, the authors find that an individual's perceived level of anonymity does increase their propensity to cyberbully, but not directly. Rather, anonymity has effects on social learning constructs that, in turn, promote cyberbullying. They close by enumerating various mechanisms by which social media platform owners could diminish users' perception of anonymity.

Ji et al. (2016) examine how to optimize a particular class of countermeasures intended to reduce the vulnerability of firms to security intrusions. A central challenge in security monitoring is to distinguish legitimate system access events from those by assailants who are doing their best to cloak their bad intentions. Rather than treating each event in isolation, companies increasingly monitor at the level of a *session*—i.e., a set of events originating from the same source. Those sessions being monitored at any given moment constitute the *hot list*. Ji et al. (2016) develop an analytical model that optimizes the size of the hot list by balancing the cost of attacks against the cost of maintaining the list itself. They complement their analytical model with a numerical simulation that ensures their approximation method is accurate and stable in a large parameter space. In addition, insights derived from their model can inform the design of socially-optimal contracts to govern outsourced session-level security monitoring activities.

2.3. Increased Interconnectedness

It's become a cliché that the Internet created a hyper-connected world, but it's no less true for that. Social networking sites, blogs, user-generated content, and peer funding (along with other sharing-economy platforms) are the latest in a long line of Internet-based technologies that allow people and organizations to connect in new ways for the purposes of communication, collaboration, and value exchange. Now physical objects increasingly interconnect due to the so-called Internet of Things (IoT). This further deepens relationships between organizations (Jernigan et al. 2016). Mobile and cloud technologies create another aspect of interconnectedness, which is that we can now *continuously* connect to our digital data, products, and knowledge work—and so can other entities, often without our knowledge.

On the positive side, interconnectedness increases the potential scale of desirable social interactions, and increases the footprint of peer production. It enables grassroots organizing and democratic social movements. Data sharing by IoT devices enables new products, industrial processes, new business models, and is associated with increased business value (Jernigan et al. 2016). Yet on the negative side, recall that a vulnerability refers to one thing's susceptibility to harm from some other thing (i.e., some entity or condition). It is inherently relational. So the more we interconnect things (people, organizations, objects) to other things, the more vulnerable they become to harm or attack from each other. Examples of vulnerabilities that are especially affected by increasing interconnectedness include security intrusions (Ji et al. 2016), cyberfraud (Lappas et al. 2016), cyberbullying (Lowry et al. 2016), and technostress (Tarafdar et al. 2015).

In this special issue, Kwon et al. (2016) examine a vulnerability very much linked to the *always-on* aspect of interconnectedness, namely, excessive dependence on social media and social games. They employ Becker and Murphy's (1988) *rational addiction* framework, which provides conceptual and modeling tools for distinguishing the extent to which people are *rational* addicts (who condition their consumption choices based on anticipated future consequences of their current behaviors) versus *myopic* addicts (who fail to recognize harmful future consequences and instead prioritize immediate gratification). Based on 13-months of panel data from thousands of smartphone users, they find that the average social app user acts in a forward-looking manner and rationally adjusts consumption over time, which suggests that rational addiction predominates in this domain. The authors also observe significant heterogeneity in the nature of addiction. Their subgroup analysis identifies several variations, such as that addictive behaviors

appear to be more myopic among older, less-educated, and higher-income groups.

Jenkins et al. (2016) also examine vulnerabilities associated with the always-on phenomenon—those caused by system-generated alerts. These alerts can provide critical information, but also increase stress and impair productivity due to dual-task interference (DTI), a cognitive phenomenon that makes it difficult for people to do two tasks at the same time. The authors depart from prior work by examining the impact of DTI on the interrupting task, rather than on the task being interrupted (the primary task). In a security context, failing to heed an alert (the interrupting task) can introduce new vulnerabilities. Functional magnetic resonance imaging (fMRI) showed that high DTI is associated with both decreased neural activity and security message disregard. Furthermore, manipulating the timing of the interruption can mitigate the effects of DTI. They argue that one way to mitigate the harm from DTI is to present warnings strategically, i.e., at times when DTI is likely to be low. Mouse cursor-tracking and psychometric measures can identify low-DTI times in security and other contexts.

2.4. Decreased Costs

One of the profound implications of digitizing something—in Negroponte’s (1996) parlance, moving from atoms to bits—is to drive the cost of perfectly replicating and distributing that thing toward zero. Some traditionally high fixed costs (e.g., high capacity digital infrastructure) are also declining. For products transitioning from atoms to bits, the effects can be substantial; a news organization, for example, no longer needs printing presses. Even within digitally native products, fixed costs are decreasing. For instance, cloud computing allows small organizations easy access to infrastructure that was previously difficult and/or expensive to set up. Individuals, startups, or units within existing firms can now begin to compete in domains that once required enormous stocks of resources to enter. As a result, many of our traditional signals of legitimacy (e.g., through investment) are no longer available; for example, a phishing email can look as legitimate as a real email.

This lowering of barriers is most pronounced in industries whose products can be completely digitized (e.g., print media), but it is also present when there is increased digital augmentation to traditional products (e.g., automobiles), or where manufacturing itself is turning digital (e.g., additive manufacturing). It is perhaps no accident that the rapid digitalization of motor vehicles preceded Tesla’s entry to the U.S. auto industry, or that established companies like GM and Ford have opened large research centers in Silicon Valley. Also, when companies develop portfolios of products, one can subsidize another. Even the traditional direction of payment can change; when data exhaust is a

byproduct, then companies may pay consumers to use their product. The result is to make companies vulnerable to competitive threats from firms that were never previously competitors. Google’s development of self-driving cars means they can threaten (or partner with) the automotive industry in unexpected ways.

The potential benefits of decreasing costs—and the lowering of related cost barriers—are many. It is generally better for products to be cheap rather than expensive, and for beneficial ideas and products to diffuse rapidly rather than slowly. The democratization of the innovation process itself—another manifestation of lowered barriers—unleashes latent energies and talents of distributed individuals, who can come together to solve problems. Worthy new voices can gain an audience, and unworthy organizations and hierarchies can be challenged on their misdeeds.

However, while digitalization lowers the costs of producing and spreading beneficial innovations and worthy ideas, it does the same for harmful innovations and dangerous or hateful ideas. Radical groups and extremists previously consigned to the fringes can gain a large social media following by fostering discontent in self-serving but societally devastating ways. Knowledge of software vulnerabilities spread rapidly among communities of attackers. “How-to” manuals for online crime, bomb-making, drug use, and other pernicious activities are freely available on the Web, as is stolen IP. Spammers and fraudsters send thousands of seemingly legitimate messages at the touch of a button. People can easily create large numbers of dummy accounts to submit fake reviews (Lappas et al. 2016), or mount undirected security attacks (Ji et al. 2016), or harass and troll at scale (see Section 3.1), or give the facade of legitimacy to fringe ideas (Manjoo 2016).

When costs rapidly decrease, the general rate of change tends to increase, which means that new kinds of harmful innovations—or even unintended consequences of beneficial ones—arise quickly compared to past eras. Of particular concern is when the rate of change in things that cause new vulnerabilities exceeds the rate of change in the mechanisms (e.g., individual behavioral patterns, organizational practices, and economic and societal structures) that we have to adapt to or counter them, or to even fully understand them. In the race between opposing innovations, harmful innovations often win (Mitra and Ransbotham 2015). Artificial intelligence (see Section 3.4) is one notable area where vulnerabilities may well spread much faster than our understanding and countermeasures.

3. Future Topics

Through these four mechanisms, new kinds of vulnerabilities—and related research opportunities—will continue to emerge and evolve as IT becomes ever more ubiquitous. Below we highlight four topics that

have attracted limited attention from information systems (IS) researchers, but that we believe are significant problems with implications at multiple levels of analysis.

3.1. Online Harassment and Incivility

Potential Problems. Social media—and other peer-production phenomena, such as open source software—empower ordinary individuals to share their ideas, talents, money, and other resources in ways that enrich society and contribute to the public good. However, the same technologies that allow the wise, well-adjusted, and kind to be heard and contribute—without first getting the sanction of established hierarchies and institutions—also empower the unwise, misguided, and cruel. In the extreme, social technologies become “a weapon of mass reputation destruction, capable of amplifying slander, bullying, and casual idiocy on a scale never before possible” (Hudson 2013).

It is not exactly news that online technologies can bring out the worst in people. A decade ago Solove (2007) wrote about the power of social media to destroy reputations, and long before that we saw “flaming” behavior in emails and on Internet message boards. Yet there is a sense that online harassment and incivility may be spiraling out of control, even to the point where it may be “ruining” the Internet (Stein 2016) or debasing political processes.

Particularly concerning is the rise of cyber mobs—swarms of people who use social technology to shame, harass, threaten, demean, or troll some targeted individual. While there is some ability to halt and reverse the harm from a single harasser, mobs are hard to fight. Cyber mobs have mounted racist and misogynous attacks on public figures, applied digital Scarlett Letters to private citizens perceived as having sinned in some way (Hudson 2013), and have even assisted individuals in carrying out vicious private vendettas (Citron 2014a).

Actions within the general category of online harassment and incivility can take many specific forms, such as:

- *Offensive speech*, which involves the communication of hateful, prurient, vulgar, or disturbing ideas that nevertheless constitute legally-protected expression;
- *Illegal harassment*, which includes real threats, defamation, and intentional infliction of emotional distress;
- *Social shaming*, which intends to hold to account a perceived norm violator;
- *Cyberbullying*, which intends to harm or demean a person in a position of relative weakness;
- *Trolling*, which intends to provoke angry or distressed responses from people as a form of amusement or for some other hidden purpose.

Trolling has, of course, long persisted as a bane of Internet forums and comments sections. While trolls can and do participate in cyber mob activities, that is not the only, or necessarily even the most common manifestation of trolling. Likewise, cyberbullies can certainly operate as individuals apart from a mob, and frequently do. By contrast, social shaming is inherently a mob activity.

Salient Research Opportunities. So far, IS researchers have primarily focused on understanding conditions that promote and enable online harassment and incivility, including personal characteristics, technological features, social learning processes, and cultural conditions (see, for example, Lowry et al. 2016). In addition to continuing this line of work, we advocate increased attention to other aspects of online harassment and incivility, as articulated in a set of research questions related to: (1) the nature and prevalence of the phenomenon, (2) consequences, and (3) mitigation strategies.

RQ1. *What is the nature and prevalence of various forms of online harassment and incivility? How is the harassment landscape evolving over time? Who is most affected?*

Quite a diverse collection of noxious behaviors constitute online harassment and incivility, ranging from mild insults to criminal acts. Furthermore, the landscape is rapidly evolving due to ongoing changes in technologies, shifts in the composition of the online population, evolution of societal attitudes, and changes to the legal environment. Thus, there is a clear need for ongoing work documenting the nature and extent of this evolving landscape. This can guide us on where to best focus our efforts at understanding enabling conditions, on the urgency of needed responses, and on whether or not progress is being made in mitigating harms.

This agenda should include the continuation of both large scale surveys, like the Pew Research Center report on Online Harassment (Duggan 2014), and focused study of particular forms such as social shaming (Solove 2007), cyber cesspools (Leiter 2010), subcultural trolling (Phillips 2015), and the newly discovered phenomenon of state-sponsored troll farms (Chen 2015).

We also advocate for investigation of new approaches based on analytics to document the extent of the phenomenon over time. For example, the same sorts of machine learning techniques that Riot Games uses to identify uncivil behavior in real time on their gaming platform (Maher 2016) could be used to benchmark the current state and ongoing evolution of civility on that site or any other.

RQ2. *What are the consequences of online harassment and incivility for individuals? What are the potential second-order effects for communities, platforms, and society?*

The harms experienced by targeted victims of the most intense forms of harassment—mental distress, loss of reputation and employment, destroyed personal relationships, and even threats to physical safety—have been thoroughly documented (see, for example, Citron 2014b). This is not to say that society at large is fully aware of these harms; however, within the research community they are fairly well-understood.

Less well-understood are potential harms caused by the background radiation of online incivility, even to people who have not been singled out for concerted attacks. Most online denizens claim that they are typically able to just ignore the incivility aimed at them (Duggan 2014), but that does not mean they remain unaffected; anxiety, depression, and social isolation could still be a result. In addition, incivility could change the way individuals process information, and not for the better. For example, Canadian researchers found that consumers who are exposed to uncivil comments regarding a scientific blog post develop different (and more extreme) views of the technology in question compared to people exposed to the exact same substantive points, but phrased in a civil way (Anderson et al. 2014).

Also of concern is the potential for toxic second-order effects. Harassment and incivility could lead to the disproportionate silencing of the minority voices most often targeted by harassers. Going further, even non-targeted segments of the population could start to withdraw from the social Internet, or they could decide to aggressively prune their social graphs to shut out the most discordant voices. In addition, continued heavy exposure to online incivility could perhaps change how people engage with the offline world, making them more extreme or entrenched in their views or how they express them, leading to a further polarization of society.

Social media companies are also vulnerable to harms (to reputational damage, loss of advertising revenue, or legislative backlash) arising from the harassment and incivility occurring on their platforms. Facebook, Twitter, and Google recently reached an agreement with Germany to remove hate speech within 24 hours after receiving notification. Measures like these could be the start of a wave of costly regulations aimed at online firms.

RQ3. Which strategies should be pursued most vigorously to mitigate the harms caused by harassment and incivility? What should platforms and technology companies be doing? What legal remedies should be taken? How can societies achieve the best balance between protection of individuals and free expression?

Citron (2014b) articulates a diverse agenda for how to mitigate the problem of online harassment and incivility, including legal reforms, more effective enforcement of existing laws, enhanced public education

efforts, and stepped up efforts by technology companies and platform owners to reign in noxious behaviors. IS scholars could help move this agenda forward by, for example:

- Devising ways to reduce technological barriers to enforcement of existing laws, e.g., by making it easier to identify attackers or preserve evidence of attacks;
- Helping legislators avoid well-intentioned but ineffective new laws (such as California's so-called Internet Erasure law, which has been criticized as technologically naïve, among other things; Lee 2014);
- Investigating how the design and governance choices that platform and community owners make can enable or constrain noxious online behaviors.

Regarding the latter, these design choices often involve difficult tradeoffs. Private businesses have the discretion to restrict expression as they see fit (since private firms are largely immune to free speech challenges); however, they need to understand the business implications and possible unintended consequences of various strategies.

IS scholars are especially well-positioned to conduct research aimed at helping platform owners navigate this difficult landscape. For example, platform owners have many levers available to adjust the level of perceived anonymity among would-be harassers (Lowry et al. 2016), but each lever has different consequences.

We can also work towards the creation of new technologies—or the assessment of existing technologies—that identify and counter noxious online behaviors in real time, perhaps in a way that is analogous to the moderately successful war against spam (Dewey 2014). For example, Jigsaw, a unit of Google, developed a prototype machine-learning tool called Conversation AI that detects apparent instances of harassment in real-time. Riot Games has experimented with tools to identify (and warn players about) toxic behaviors during game play on their League of Legends platform (Maher 2016).

To sum up, IS scholars can do much to advance research on the conditions enabling online harassment, its effects, and mitigation strategies. While cyberspace is inherently more difficult to govern and police than physical spaces, it also offers much greater opportunity for technological solutions.

3.2. Exacerbation of Inequality

Potential Problems. Digital technologies increasingly drive productivity improvement and innovation in modern economies. However, at the same time, these technologies may reinforce a less salutatory economic trend, which is an increase in economic inequality. Since the 1970s, economic inequality has risen sharply in most industrialized countries (Alvarado et al. 2013, Atkinson et al. 2011, Piketty 2014). In the

U.S., the top 1% now hold 35% of the wealth (Wolff 2010) and 23% of income (Saez 2013).

Extreme economic inequality can lead to a number of harms, including inequality of opportunity, dampened of consumer demand, inadequate investment in public infrastructure, corruption of political institutions, and even social breakdown or violent revolutions (Acemoglu and Robinson 2012, Chapter 12; Stiglitz 2014; Brynjolfsson and McAfee 2014, Chapter 11). Furthermore, across a wide variety of countries and time periods, higher inequality has been associated with lower rates of economic growth (Ostry et al. 2014).

So, it seems clear that extremes of inequality can make individuals and societies more vulnerable. Yet, in what sense should we view this as a *digital* vulnerability, i.e., one that stems from the growing ubiquity of digital technologies? Are digital technologies a big part of the rising inequality story, or just a footnote? If they are more than a footnote, how can we mitigate their contributions to growing inequality?

Drivers of Digital Inequality. Many of the factors proposed to explain growing inequality have no clear connection to digitalization, such as weakening of the labor movement, decreases in tax code progressivity, and changes to norms about executive compensation (Stiglitz 2012, Piketty 2014). However, others do relate to digitalization (Brynjolfsson and McAfee 2014, Chapter 10; Brynjolfsson et al. 2014). We highlight seven of these factors below.

Globalization: Globalization and offshoring promote inequality in wealthier countries by, among other things, putting downward pressure on wages (Stiglitz 2014). IT is, of course, increasingly essential to coordinating global supply chains.

Skill-biased Technical Change: Skill-biased technical change is when technological innovations disproportionately increase the value of skilled workers or decrease the value of unskilled workers, for example, by using automation to replace low-skilled workers, or using decision technologies to augment high-skilled workers (Autor et al. 2008).

Persistent Technological Unemployment: Ordinarily, job-destroying technologies only boost unemployment in the short run; the economy adjusts, and labor moves to other productive uses. However, this logic assumes that the ongoing rate of adjustment (e.g., through worker retraining) will be fast enough to offset the ongoing rate of technological displacement (Brynjolfsson and McAfee 2011, Chapter 3).

Power Law Performance Distributions: If it is true that job performance increasingly follows a power law (O’Boyle and Aguinis 2012)—or even if compensation is increasingly allowed to reflect existing distributions that already are, or are perceived to be, power law—then the result will be to increase wage inequality.

Because of digitalization, it is plausible that innovators and others in the “creative class” will increasingly experience a power law distribution of compensation, one reason being that digitization allows creative outputs to be more easily replicated (at low cost) over a large audience (Brynjolfsson et al. 2014).

Network Effects: A well-established feature of network markets is their propensity to exhibit winner-take-all outcomes, in which one or a few suppliers dominate a market and thereby earn outsized rewards. Increasing digitalization means that an increasing number of markets will be subject to network effects, which could further contribute to rising inequality.

Online Labor Markets: One manifestation of the digital economy is the rise of online labor markets such as Upwork, Amazon’s Mechanical Turk, and Uber. While these markets currently account for only a small proportion of the rapid rise of alternative work arrangements comprising the so-called “gig,” “freelance,” or “1099” economy, its share is growing rapidly (Katz and Krueger 2016). These markets can increase inequality (especially in wealthy countries) through several mechanisms, such as by increasing opportunities for global labor arbitrage, by shifting risk from employers to workers, and by increasing the salience of the “Matthew Effect”) (see below).

The Digital Divide: The digital divide refers to the large disparity between the rich and poor in their access to, or ability to exploit, digital technologies (Wei et al. 2011). Digital technologies increasingly mediate our access to products, services, education, and job opportunities. If the level of disparity were to grow—or even if it were to stay the same while the economic importance of the disparity were to grow due to the increasing ubiquity of digitalization—then it will exacerbate inequality.

Salient Research Opportunities. Two broad avenues for future work on the link between digital ubiquity and rising inequality are particularly salient for IS scholars, as embodied in the two research questions we elaborate below.

RQ1. *By what theoretical mechanisms does increasing digital ubiquity contribute to economic inequality? Which are the most important contributors?*

As a first step for future research, we call for increased attention to potential IT-related drivers of inequality—such as the seven enumerated above—to better understand the specific mechanisms by which they contribute to inequality. Some of these drivers (e.g., skill-biased technical change, technological unemployment) have received considerable attention already.

However, for other drivers, our understanding is just beginning. One of these is the apparent trend towards power law performance distributions. One job that seems to follow a power law distribution is computer programming, where studies have documented

order-of-magnitude variations in productivity, even among employees with similar experience (Valett and McGarry 1989). A large skewing of performance also seems likely in the increasing number of other positions where the focus is on technological innovation and/or creative problem solving. Yet, care is needed when measuring job performance. For example, it appears that some of O’Boyle and Aguinis’ (2012) empirical findings were driven more by artifacts of measurement than underlying performance distributions (Beck et al. 2014). That said, these same artifacts may be increasingly appearing in compensation schemes, for good or ill (Bock 2015, Chapter 10). For example, one of seven artifacts identified by Beck et al. (2014) is to only give people credit for extreme performances, such as to only count the number of Emmy’s an actor has won—or the number of “A” papers a scholar has published.

Another trend worthy of attention is the growth of online labor markets (Kokkodis and Ipeirotis 2015). These markets cause a shift in risks (and associated economic burdens) from organizations to individuals, and also enable workers from low wage nations to compete for work with those from high wage nations (resulting in loss of work and/or a lowering of pay rates for the latter). The risk shift is seen most clearly when work is done “on spec” (i.e., many people submit work products but only one gets paid) but is actually a feature of any platform that allows employers to replace full-time employees with an on-demand labor force. More interestingly, the “Matthew Effect”—wherein a person’s initial success in some endeavor increases their visibility and access to resources, thus begetting more success—may be especially pronounced in online labor markets. Such markets typically make visible an individual’s detailed work history, and may even publish rankings of workers, either of which should promote the Matthew Effect. In fact, digital ranking and filtering tools, which are increasingly used in job candidate searches and hiring in general, may disproportionately advantage those who have already achieved certain credentials, and could therefore increase the trend toward power law performance distributions (Brynjolfsson and McAfee 2014, Chapter 10).

RQ2. How can we mitigate the contributions of digital ubiquity to growing economic inequality, or use digital technology as a tool for mitigation itself?

As already noted, many important drivers of inequality are unrelated to technology, and for these there are some known remedies²—which is not to say

that any of them will be easy to get adopted in the U.S. or other countries with a similar political and social environment. However, our focus here is specifically on (1) how inequality drivers that are especially caused or exacerbated by digital technology can be mitigated, and (2) how digital technology itself can be used as a potential tool for mitigation.

Regarding the former, technological unemployment could be countered in several ways, including: wage insurance (Ghilarducci 2016), the guaranteed basic or minimum income, negative taxation (Friedman 2009), increased investment in education and training, and encouragement of entrepreneurship and innovation (Brynjolfsson and McAfee 2014, Chapters 13 and 14).

Regarding the latter question, technology itself could mitigate growing inequality in several ways. For example, policies could promote investment in “resource-saving technological change” by making firms pay higher fees for the environmental impacts of their production (Stiglitz 2012, Chapter 3). Technology can also increase the quality and availability of education and training (e.g., through Massive Open Online Courses (MOOCs)), and can help employers better match candidates to job opportunities (Brynjolfsson and McAfee 2014, Chapter 13). Yet, as always, unanticipated consequences lurk within these strategies. If cash-strapped states seek to shift poorer residents from high cost state universities to low cost MOOCs, the net result could be to increase inequality. Better job matching algorithms could encourage additional hiring, but as already noted, they could also reinforce the Matthew Effect.

In sum, much interesting work remains to be done at the nexus of technology and inequality, especially with regard to policy, which is a domain some have argued should receive more attention from IS scholars in any case (Lucas et al. 2013).

3.3. Industrial Internet of Things

Potential Problems. The Industrial Internet of Things (IIoT) focuses, as the name suggests, on industrial (rather than consumer) uses of IoT. As such, the IIoT is directed towards large-scale industrial applications, primarily in energy, transportation, manufacturing, agriculture, and healthcare.³

The technical side of IIoT is about enhancing industrial machines with digital sensors, actuators, and local intelligence, and then connecting them to each other and remote computers over wireless networks. By contrast, the strategic side of IIoT aims at understanding

² Stiglitz (2012, Chapter 10) lists a number of such remedies, including adoption of regulatory changes to discourage rent seeking, changes to rules for executive compensation, reform of the tax code and broader tax policy, and increases in public investment in education, health, and infrastructure.

³ “Industry 4.0” is a related term used in Europe that focuses more narrowly on manufacturing. The name stems from the assertion that industry has progressed through four revolutions. The first three centered on mechanization, mass production, and computer automation, respectively, while the fourth, Industry 4.0, centers on integrated cyber-physical systems.

how to combine connected machines, data analytics, and new operating and business models to generate value.

The hopes for value generation are high, to say the least. Accenture, McKinsey, and General Electric (which has staked much on its “Industrial Internet” strategy) each project trillions of dollars in cumulative value creation potential over the next 15 years. A typical simplifying assumption underlying such projections is that IIoT could result in about a 1% improvement of industrial productivity worldwide by 2030.

Salient Research Opportunities. If such projections might seem optimistic for any number of reasons, it is still not hard to imagine substantial improvements in favorable circumstances. It is also not hard to imagine sources of vulnerability when everything contains a computer...and those computers interconnect...and we lack physical control of the devices...and the devices control objects in the physical world...and they operate autonomously. Yet for IIoT to be a distinctly important research opportunity, there needs to be ways in which IIoT challenges some of our prior understandings. If it is only “more computers,” then that may not warrant distinct research efforts. We identify four ways in which IIoT may be more than just “more.” These pertain to: (1) implications of measurement, (2) control of physical objects, (3) organizational relationships, and (4) embedded complexity.

RQ1. How will better measurement of real-world activities affect management? How do biases change when organizations know more?

Most IIoT applications involve monitoring and measuring. Their precise intent is to gather more data about what is happening in the real world. Many attributes that are currently only surmised will have evidence. While real-time streams of data will certainly affect IT infrastructure and governance, the implications of detailed measurement data are far less clear. Managers may focus on what is measurable instead of what is important.

Furthermore, measurement, even though improved, will still be imperfect. Understanding the measures and results from them will require organizations to develop more skills around the interpretation of analytical results and the management of analytical processes that involve complex, multidimensional data from myriad sources. Organizations may face increasing vulnerability to many forms of biases, all exacerbated by volumes of detailed measurements.

What’s more, with increased ability to measure, organizations may face increasing responsibility to act. Measurements from IIoT devices will provide considerable ammunition for second guessing. Given a real-time stream of operational data, what responsibility does an organization have to prevent an industrial

accident? Opportunities may arise to design economic mechanisms that use this data to better align incentives and consequences.

RQ2. How do we adapt security and countermeasures for a blended virtual and physical world?

IIoT devices can go far beyond measurement only. No longer passive observers, the devices will work with autonomy. This builds on a long precedent; for example, thermostats turn heating systems on and off to maintain a desired temperature. Yet as IIoT devices control more complex and more dangerous physical components, the stakes can be far greater than a room that is too hot or too cold.

As a result, all of the information security issues in electronic systems are harder and more consequential. Absence of physical controls adds to the already complex information security problem. Autonomous actions in the physical world add to the repercussions of security shortcomings.

RQ3. Can we align incentives so that the organizations can interact with other organizations they depend on but are also at risk from?

IIoT devices not only connect to other devices, they create connections between organizations (Jernigan et al. 2016). Some are straightforward. An equipment manufacturer may depend on the data from another organization, such as a customer using the equipment, to provide data required to understand wear and maintenance. Yet even that straightforward example may introduce dependencies on other organizations, such as telecommunications providers. The challenges associated with managing robust interconnected systems are fraught with disincentives and externalities—in the best case. For example, what are equitable mechanisms for prioritizing data transfer between IIoT devices when resources are constrained?

Bad actors accentuate these difficulties. The stakes for trust between systems are high when tight interconnections between organizations create shared, systemic risk. The potential for collateral damage looms large when the integrity of some devices are compromised, setting off a chain reaction. The allegiance of these distributed devices may be fickle as attackers conscript them for nefarious purposes that may affect us all, such as the recent distributed denial of service attack on domain name service provider Dyn. Many of these concerns go beyond technical. For example, if blockchains are able to provide distributed integrity, how do organizations standardize on protocols on distributed devices?

RQ4. How do non-software organizations navigate the complexities of embedded software?

Despite years of experience, software development companies still struggle with vulnerabilities. Yet, the IIoT requires even non-software companies to become proficient at not only their traditional products but also

the software components that will become embedded in them. This is a tall order.

Initial forays into IIoT are likely to be filled with mistakes due to inexperience, thus creating vulnerabilities that nefarious actors can exploit. Prevention requires knowledge that is currently in short supply. Until IIoT components are more mature, organizations face not only the challenges of developing the IIoT devices but also difficulties in managing the inevitable consequences of weaknesses.

3.4. Algorithmic Ethics and Algorithmic Bias

Potential Problems. We increasingly endow devices and systems with intelligent algorithms that guide their operations in autonomous ways. Algorithms can improve operational performance and free up human attention for other tasks, but because these devices can do direct harm (including physical) to people, they also raise thorny new ethical questions. In some cases, stakeholders must decide what level of autonomy to give an algorithm, such as whether drones should make unilateral “kill decisions.” In other cases, as in self-driving cars, autonomy is assumed, and so the question becomes which principles should guide this autonomy? Should a self-driving car swerve around an obstacle to protect the driver even if this would put the lives of pedestrians at risk (Bonnefon et al. 2016)?

In addition, the increasing ubiquity of IT creates vast new troves of individual data and ever more clever matching algorithms, which raises the specter of *algorithmic bias*—computer algorithms that have the effect (intentional or unintentional) of unfairly penalizing certain groups or diminishing their access to some product or opportunity. On the plus side, algorithms can be a great boon in terms of matching marketing messages to those who will be most receptive; consumer products to those who most want and need them; entertainment products to those who will most enjoy them; insurance products and medical treatment to people with the right risk and health profiles; and job and educational opportunities to those who will be the strongest contributors. Yet inherent in the concept matching is the idea of discernment, and so it should come as no surprise that scholars are increasingly concerned about the potential of IT to promote the ugly side of discernment, i.e., bias and discrimination (Barocas and Selbst 2016). The same technology that empowers us to make the distinctions necessary to sort people into categories appropriately can be used to sort people inappropriately, or to deny them access to products and opportunities that they deserve.

Salient Research Opportunities. We advocate attention going forward to both algorithmic ethics and the related issue of algorithmic bias, as encapsulated in two sets of research questions presented below.

RQ1. How do we ensure appropriate human dominion over and accountability for the actions of intelligent autonomous machines? What is a legal and ethical framework for assigning responsibility for the harmful actions of intelligent algorithms?

Interest among researchers and industry groups about the ethical implications of intelligent algorithms is growing. Last year, Russell et al. (2015b), and a companion article in *AI Magazine* (Russell et al. 2015a) advocated a set of research priorities for “robust and beneficial Artificial Intelligence (AI).” More recently, five leading technology companies announced an initiative to create a standard of ethics to govern the development and deployment of AI technologies (Markoff 2016). In yet another industry-led effort, several Silicon Valley entrepreneurs committed \$1 billion to fund a nonprofit called OpenAI (Thornhill 2015), whose stated mission is to “build safe AI, and ensure AI’s benefits are as widely and evenly distributed as possible.”

These efforts are being made against a backdrop of public alarm about whether the capabilities of intelligent algorithms will race ahead of our wisdom about how best to control them—or even our ability to control them (*The Economist* 2016). Some of these concerns may be premature or overblown, but it is certainly true that algorithms will continue to encroach on decisions once reserved for humans, including ones that provoke ethical quandaries. If we are not proactive in anticipating and mitigating potential harms posed by intelligent algorithms (and the related ethical dilemmas they may pose) then we may face an elevated risk that an AI Pandora’s box really might be opened, or if not that, then fears of this possibility could provoke a public backlash or ill-advised attempts at regulation.

Since answers to the research questions posed above lie squarely at the intersection of technology, business, and ethics, we believe IS researchers are especially well-positioned to contribute. It would be unfortunate if this work were left entirely in the hands of those industry groups or individuals whose natural self-interest might leave them especially prone to a *pro-innovation bias* with respect to these technologies.

RQ2. How can we know when algorithms are unfairly biased? How can we ensure their decisions will be reviewable and accountable? What are organizational best practices to avoid or mitigate potential harms from algorithmic bias?

Algorithmic bias—also referred to as *data discrimination*—is becoming a matter of significant public concern, as reflected in news articles (Kirchner 2015, Miller 2015) and a recent Obama administration report on this topic (Munoz et al. 2016, p. 25). The latter report defines discrimination “in a very broad sense to refer to outsized harmful impacts—whether intended or otherwise—that the design, implementation, and

utilization of algorithmic systems can have on discrete communities and other groups that share certain characteristics . . . Some instances of discrimination may be unintentional and even unforeseen. Others may be the result of a deliberate policy decision to concentrate services or assistance on those who are most in need. Still others may create adverse consequences for particular populations that create or exacerbate inequality of opportunity for those already at a disadvantage.”

This definition highlights a central challenge of this emerging stream: to understand what mechanisms can generate algorithmic bias, and how to tell when potential bias is in fact present. Looking beyond those (no doubt) rare cases where designers had discriminatory intent (and this can be discovered), scholars working in the emerging field of *algorithmic accountability* identify some more subtle mechanisms. For example, data could be flawed in a way that causes algorithms to reflect biases of prior decision makers, or data could be valid, but still reflect historical biases that exist in the broader society (Barocas and Selbst 2016).

As a result, it will be difficult to know if discrimination really is present—and when it is, what to do about it. A controversial new report illustrates this conundrum well. Broward County, Florida, predicts recidivism among felons using an algorithm. ProPublica found that blacks were almost twice as likely as whites to be false positives, i.e., to be labeled by the algorithm as a “higher risk” to reoffend, but not actually reoffend (Angwin et al. 2016). The algorithm also skewed with regard to false negatives; i.e., whites classified as “lower risk” were much more likely than blacks to reoffend.

The algorithm developer (Northpointe, Inc.) pushed back, pointing out that the algorithm exhibits similar levels of *sensitivity* (true positive rate) and *specificity* (true negative rate) for blacks and whites (Dieterich et al. 2016). Northpointe argued that differences in sensitivity and specificity across racial groups would be bias, not differences in the model error rate. They pointed out that the higher model error rates among blacks is a statistical artifact resulting (in part) from the higher base rate of recidivism in this group. Yet there is intuitive appeal in ProPublica’s contention that race-correlated prediction errors are a sign of a problem. When judging *disparate impact*, a legal principle applied by some U.S. courts in discrimination cases brought under Title VII of the Civil Rights Act (Kirchner 2015), which standard should be used? Irrespective of the decisions, long buried biases and trade-offs will come to light when codified.

It will take a combination of skills (technical, statistical, legal, social scientific) to identify biased algorithms (Pasquale 2015). Going forward, we think there is a great opportunity for IS scholars to contribute to these sorts of interdisciplinary efforts.

4. Conclusion

While we selected four topics to feature as promising areas of research (see Table 1 for a summary), clearly there are numerous others. Research tends to gravitate towards positive aspects of innovation while negative aspects receive less attention (Rogers 2003); this creates an opportunity for IS scholars interested in digital vulnerabilities. Many of the established regulatory mechanisms (such as laws, inspections) with a mature history in the physical world struggle in the digital world, where industry and national boundaries are weak and negative externalities are strong. Some additional kinds of vulnerability worthy of study that are emerging from our increasingly digital world include those arising from:

- **Economic effects of AI on individuals, organizations, and nations:** What are the theoretical mechanisms by which AI and robotics could lead to sustained mass unemployment? Which populations would be most affected? What countermeasures could moderate these job losses?
- **Omnipresence:** Information technologies are pervasive. How does the constant, permeating use of technology (such as mobile devices, social media, IoT) affect people (e.g., technostress, surveillance, electronic discovery, negative word of mouth)? How can organizations connect with consumers without attention becoming a tragedy of the commons? How can we mitigate real world effects of vulnerabilities within our digital personas (Schultze and Mason 2012)?
- **Multipurpose technologies:** Few technologies are purely positive or negative; there is a vast middle ground with relative assessments of positive or negative highly dependent on perspective. How can we benefit from the efficiencies and transparencies of distributed ledgers (such as blockchain) even beyond currencies without enabling nefarious uses?
- **Systemic transitions:** Digitization affects practically every segment of society, often supplanting long established norms and systems. How do we transition from old systems to new—such as in elections, transportation, education—without provoking disastrous unintended consequences? How do we mitigate the systemic versus idiosyncratic risks from technology dependence? When more becomes transparent, how do we counter unproductive gaming of systems?

While data is increasingly available, both in quantity and quality, to support innovative research in these domains, we need to take care to study problems because they are important, not just because data is convenient. How can IS research benefit society by helping maximize benefits and create value from advances in IT while minimizing drawbacks? Beyond just identifying vulnerabilities and understanding the mechanisms that cause them, IS research can lead by

Table 1 Four Directions for Research on Digital Vulnerabilities

Research topic	Potential harms	Link to mechanisms	Key research questions
Online harassment and incivility	<p>For victims—Mental anguish, loss of reputation, loss of work, being forced off social media</p> <p>For platforms—Damage to brand image, escalation of monitoring costs, potential imposition of onerous new regulation</p> <p>For society—Silencing of minority voices, “ruining of the Internet,” fostering discontent for self-serving purposes</p>	<p><i>Visibility</i>: Provides would-be harassers with personal information about victims. The acts of harassment itself becomes visible to a much larger audience. This attracts people who want an audience (e.g., trolls) and multiplies the harm to the victim</p> <p><i>Cloaking</i>: Harassers and others who engage in incivility often disguise their true motives and/or hide behind anonymity</p> <p><i>Interconnectedness</i>: Growing interconnectedness makes the reputational effects of harassment especially pernicious. It also increases the cost of opting out to avoid further harassment</p> <p><i>Low costs</i>: Harassment no longer requires a physical presence, or even waiting for a phone to connect. One person can easily harass scores of people. New harassment tactics emerge and evolve rapidly (e.g., cybermobs, doxxing, malicious impersonation)</p>	<p>How prevalent are various forms of online harassment? What are the key enablers? Who is engaging in it the most and why?</p> <p>Who is most subject to harassment? Why? What are the harms?</p> <p>Which mitigation strategies are most effective at protecting harassment victims? What should platforms and governments be doing about harassment and incivility? How can we achieve the best balance between protection and free expression?</p>
Exacerbation of inequality	<p>Unequal opportunity, diminished economic growth, social unrest, corruption of political institutions, economic breakdown, violent insurrection</p>	<p><i>Visibility</i>: Reputations and some kinds of performance become more visible. This can reinforce power law performance/compensation schemes</p> <p><i>Cloaking</i>: Those at the favorable end of the inequality equation have incentives to cloak key ideas and skills to preclude others from competing effectively</p> <p><i>Interconnectedness</i>: Promotes network effects, which can lead to winner-take-all outcomes</p> <p><i>Low costs</i>: AI robots (and other software technologies that can substitute for human labor) can be replicated at a near zero cost. Pace of change may be too fast for normal economic forces to reallocate labor in a way that avoids persistent unemployment</p>	<p>What are the precise mechanisms by which increased digitalization contributes to inequality?</p> <p>What is the relative contribution of these mechanisms to the overall rise in inequality?</p> <p>What mitigation strategies can be used to moderate the effects of these mechanisms?</p>
Industrial Internet of Things	<p>Physical objects and industrial systems increasingly under attack</p> <p>Attacks have increasingly large ripple effects</p> <p>The fates of firms engaging in “outcomes-based” contracting become entwined in ways that can lead to a cascade of failures</p>	<p><i>Visibility</i>: IIoT is fundamentally about making physical things visible in new ways</p> <p><i>Cloaking</i>: Data thought to be from an organization’s devices might not be</p> <p><i>Interconnectedness</i>: IIoT all about connecting up physical things in new ways, and creating new connections between organizations.</p> <p><i>Low costs</i>: IIoT is enabled by a low cost of adding sensors and software to objects</p>	<p>How will better measurement of real-world activities affect management? How do biases change when organizations know more?</p> <p>How do we adapt security and countermeasures for a blended virtual and physical world?</p> <p>Can we align incentives so that the organizations can interact with other organizations they depend on but are also at risk from?</p> <p>How do non-software organizations navigate the complexities of embedded software?</p>
Algorithmic ethics and algorithmic bias	<p>Autonomous machines make unethical choices</p> <p>Opportunities (or sanctions) are presented to people in a discriminatory way due to algorithmic biases</p>	<p><i>Visibility</i>: Increased visibility of personal data increases the domains in which individuals could be subject to algorithmic bias</p> <p><i>Cloaking</i>: One feature of intelligent algorithms is opaqueness in how they reach decisions, and opaqueness about who is responsible</p> <p><i>Interconnectedness</i>: Algorithms that interact across different domains, for example, health and financial, can amplify and create new vulnerabilities that do not arise in individual domains</p> <p><i>Low costs</i>: Cost of deploying intelligent algorithms is dropping rapidly, which expands the domain of potential bias</p>	<p>How do we ensure appropriate human dominion over and accountability for the actions of intelligent autonomous machines? What is a legal and ethical framework for assigning responsibility for the harmful actions of intelligent algorithms?</p> <p>How can we know when algorithms are biased? How do we make their decisions reviewable and accountable? What are organizational best practices to avoid or mitigate potential harms from algorithmic bias?</p>

developing new managerial wisdom for better managing organizations, societies, and our personal lives in light of these vulnerabilities.

Appendix. Special Issue Process

On August 27, 2014, the special issue editors and *ISR* editor-in-chief Ritu Agarwal invited researchers to focus on the dark side of IT through a call for papers for a special issue. Interested researchers could submit paper ideas for early reaction from the senior editors by January 4, 2015; full papers were due by March 1, 2015. The editorial team screened all submitted papers and moved papers deemed to have a reasonable chance of acceptance in an accelerated time frame into the review process. After an initial round of review, the editors invited authors of papers still under consideration to present at a workshop on September 19, 2015, at Boston College. After further review rounds, the final decision for the special issue was made in August 2016.

The special issue editors appreciate the assistance of *ISR* editor-in-chief Ritu Agarwal and the editorial board. Board members were

- Ahmed Abbasi (Virginia)
- Alessandro Acquisti (CMU)
- Terrence August (UCSD)
- France Belanger (Virginia Tech)
- Huseyin Cavusoglu (UT Dallas)
- Ram Chellappa (Emory)
- Elizabeth Davidson (Hawaii)
- Debabrata Dey (Washington)
- Kai Lung Hui (HKUST)
- Eric Johnson (Vanderbilt)
- Karthik Kannan (Purdue)
- Paul Lowry (Hong Kong)
- James Marsden (Connecticut)
- Sabyasachi Mitra (Georgia Tech)
- Vijay Mookerjee (UT Dallas)
- Tyler Moore (SMU)
- Eric Overby (Georgia Tech)
- Rema Padman (CMU)
- Srinivasan Raghunathan (UT Dallas)
- Sasha Romanosky (RAND)
- Matti Rossi (Aalto)
- Larry Sanders (SUNY–Buffalo)
- Raghu Santanam (Arizona State)
- Ulrike Schultze (SMU)
- Olivia Sheng (Utah)
- Param Singh (CMU)
- Mikko Siponen (Oulu)
- Carsten Sørensen (LSE)
- Catherine Tucker (MIT)
- Merrill Warkentin (Mississippi State)

References

- Acemoglu D, Robinson J (2012) *Why Nations Fail: The Origins of Power, Prosperity, and Poverty* (Crown Business, New York).
- Akers RL (2011) *Social Learning and Social Structure: A General Theory of Crime and Deviance* (Northeastern University Press, Boston).
- Alvarado F, Atkinson AB, Piketty T, Saez E (2013) The top 1 percent in international and historical perspective. *J. Econom. Perspect.* 27(3):3–20.
- Anderson AA, Brossard D, Scheufele DA, Xenos MA, Ladwig P (2014) The “nasty effect”: Online incivility and risk perceptions of emerging technologies. *J. Computer-Mediated Comm.* 19(3):373–387.
- Angwin J, Larson J, Mattu S, Kirchner L (2016) Machine bias. *ProPublica* (May 23), <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing/>.
- Atkinson AB, Piketty T, Saez E (2011) Top incomes in the long run of history. *J. Econom. Literature* 49(1):3–71.
- Autor DH, Katz LF, Kearney MS (2008) Trends in U.S. wage inequality: Revising the revisionists. *Rev. Econom. Statist.* 90(2):300–323.
- Barocas S, Selbst AD (2016) Big data’s disparate impact. *California Law Rev.* 104(3):671–732.
- Beck JW, Beatty AS, Sackett PR (2014) On the distribution of job performance: The role of measurement characteristics in observed departures from normality. *Personnel Psych.* 67(3):531–566.
- Becker GS, Murphy KM (1988) A theory of rational addiction. *J. Political Econom.* 96(4):675–700.
- Bock L (2015) *Work Rules!: Insights from Inside Google That Will Transform How You Live and Lead* (Twelve, New York).
- Bonnefon JE, Shariff A, Rahwan I (2016) The social dilemma of autonomous vehicles. *Science* 352(6293):1573–1576.
- Brynjolfsson E, McAfee A (2011) *Race Against the Machine: How the Digital Revolution is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy* (Digital Frontier Press, Lexington, MA).
- Brynjolfsson E, McAfee A (2014) *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies* (W. W. Norton & Company, London).
- Brynjolfsson E, McAfee A, Spence M (2014) New world order: Labor, capital, and ideas in the power law economy. *Foreign Affairs* 93(4):44–53.
- Cavusoglu H, Phan TQ, Cavusoglu H, Airolidi ED (2016) Assessing the impact of granular privacy controls on content sharing and disclosure on Facebook. *Inform. Systems Res.* 27(4):848–879.
- Chen A (2015) The agency. *New York Times Magazine* (June 2), <http://www.nytimes.com/2015/06/07/magazine/the-agency.html>.
- Citron DK (2014a) How cyber mobs and trolls have ruined the Internet—and destroyed lives. *Newsweek* (September 19), <http://www.newsweek.com/internet-and-golden-age-bully-271800/>.
- Citron DK (2014b) *Hate Crimes in Cyberspace* (Harvard University Press, Cambridge, MA).
- Dewey C (2014) Is spam free speech? *Washington Post* (July 31), <https://www.washingtonpost.com/news/the-intersect/wp/2014/07/31/is-spam-free-speech/>.
- Dieterich W, Mendoza C, Brennan T (2016) COMPAS risk scales: Demonstrating accuracy equity and predictive parity performance of the COMPAS risk scales in Broward County. Northpointe Inc. Report. http://go.volarisgroup.com/rs/430-MBX-989/images/ProPublica_Commentary_Final_070616.pdf.
- Duggan M (2014) Online harassment. Pew Research Center, Washington, DC.
- Friedman M (2009) *Capitalism and Freedom* (University of Chicago Press, Chicago).
- Ghilarducci T (2016) What is this “wage insurance” Obama’s talking about? *The Atlantic* (January 14), <http://www.theatlantic.com/business/archive/2016/01/wage-insurance-state-union/424167/>.
- Hudson L (2013) Why you should think twice before shaming anyone on social media. *Wired Magazine* 21(8) https://www.wired.com/2013/07/ap_argshaming/.
- Jenkins JL, Anderson BB, Vance A, Kirwan CB, Eargle D (2016) More harm than good? How messages that interrupt can make us vulnerable. *Inform. Systems Res.* 27(4):880–896.
- Jernigan S, Ransbotham S, Kiron D (2016) Data sharing and analytics drive success with IOT. *Sloan Management Rev.* (September 8), <http://sloanreview.mit.edu/projects/data-sharing-and-analytics-drive-success-with-internet-of-things/>.
- Ji Y, Kumar S, Mookerjee V (2016) When being hot is not cool: Managing hot lists in information security. *Inform. Systems Res.* 27(4):897–918.

- Katz LF, Krueger AB (2016) The rise and nature of alternative work arrangements in the United States, 1995–2015. Working paper, Harvard University, Cambridge, MA.
- Kirchner L (2015) When discrimination is baked into algorithms. *The Atlantic* (September 6), <http://www.theatlantic.com/business/archive/2015/09/discrimination-algorithms-disparate-impact/403969/>.
- Kokkodis M, Ipeirotis PG (2015) Reputation transferability in online labor markets. *Management Sci.* 62(6):1687–1706.
- Kwon EK, So H, Han SP, Oh W (2016) Excessive dependence on mobile social apps: A rational addiction perspective. *Inform. Systems Res.* 27(4):919–939.
- Lappas T, Sabnis G, Valkanas G (2016) The impact of fake reviews on online visibility: A vulnerability assessment of the hotel industry. *Inform. Systems Res.* 27(4):940–961.
- Lee J (2014) SB 568: Does California’s online eraser button protect the privacy of minors? *U.C. Davis Law Rev.* 48(3):1173–1206.
- Leiter B (2010) Cleaning cyber-cesspools: Google and free speech. Levmore S, Nussbaum MC, eds. *The Offensive Internet: Speech, Privacy, and Reputation* (Harvard University Press, Cambridge, MA), 155–173.
- Lowry PB, Zhang J, Wang C, Siponen M (2016) Why do adults engage in cyberbullying on social media? An integration of online disinhibition and deindividuation effects with the social structure and social learning model. *Inform. Systems Res.* 27(4):962–986.
- Lucas HC Jr, Agarwal R, Clemons EK, El Sawy OA, Weber BW (2013) Impactful research on transformational information technology: An opportunity to inform new audiences. *MIS Quart.* 37(2):371–382.
- Maher B (2016) Can a video game company tame toxic behaviour? *Nature* (March 30), http://www.nature.com/news/can-a-video-game-company-tame-toxic-behaviour-1.19647#.
- Manjoo F (2016) Social media’s globe-shaking power. *New York Times* (November 16), <http://www.nytimes.com/2016/11/17/technology/social-medias-globe-shaking-power.html>.
- Markoff J (2016) How tech giants are devising real ethics for AI. *New York Times* (September 1), http://www.nytimes.com/2016/09/02/technology/artificial-intelligence-ethics.html?_r=0.
- Miller CC (2015) When algorithms discriminate. *New York Times* (July 9), <http://www.nytimes.com/2015/07/10/upshot/when-algorithms-discriminate.html>.
- Mitra S, Ransbotham S (2015) The effects of vulnerability disclosure policy on the diffusion of security attacks. *Inform. Systems Res.* 26(3):565–584.
- Munoz C, Smith M, Patil D (2016) Big data: A report on algorithmic systems, opportunity, and civil rights. *Executive Office of the President*, https://www.whitehouse.gov/sites/default/files/microsites/ostp/2016_0504_data_discrimination.pdf.
- Negroponte N (1996) *Being Digital* (Vintage Books, New York).
- O’Boyle E, Aguinis H (2012) The best and the rest: Revisiting the norm of normality of individual performance. *Personnel Psych.* 65(1):79–119.
- Ostry JD, Berg AM, Tsangarides CG (2014) Redistribution, inequality, and growth. *Internat. Monetary Fund*, <http://www.imf.org/external/pubs/ft/sdn/2014/sdn1402.pdf>.
- Overby E, Slaughter S, Konsynski B (2010) The design, use, and consequences of virtual processes. *Inform. Systems Res.* 21(4):700–710.
- Pasquale F (2016) Digital star chamber. *Aeon*, <https://aeon.co/essays/judge-jury-and-executioner-the-unaccountable-algorithm>.
- Phillips W (2015) *This Is Why We Can’t Have Nice Things: Mapping the Relationship Between Online Trolling and Mainstream Culture* (MIT Press, Cambridge, MA).
- Piketty T (2014) *Capital in the Twenty-First Century* (Harvard University Press, Cambridge, MA).
- Rogers E (2003) *The Diffusion of Innovations* (Free Press, New York).
- Russell S, Dewey D, Tegmark M (2015a) Research priorities for robust and beneficial artificial intelligence. *AI Magazine* 34(4):105–114.
- Russell S, Dietterich T, Horvitz E, et al. (2015b) An open letter: Research priorities for robust and beneficial artificial intelligence. *Future of Life Institute*, <http://futureoflife.org/ai-open-letter/>.
- Saez E (2013) Striking it richer: The evolution of top incomes in the United States (updated with 2012 preliminary estimates). University of California, Department of Economics, Berkeley. <http://elsa.berkeley.edu/~saez/saez-UStopincomes-2012.pdf>.
- Schultze U, Mason RO (2012) Studying cyborgs: Re-examining Internet studies as human subjects research. *J. Inform. Tech.* 27(4):301–312.
- Solove DJ (2007) *The Future of Reputation* (Yale University Press, New Haven, CT).
- Stein J (2016) How trolls are ruining the Internet. *Time Magazine* (August 18), <http://time.com/4457110/internet-trolls/>.
- Stiglitz JE (2012) *The Price of Inequality: How Today’s Divided Society Endangers Our Future* (WW Norton & Company, New York).
- Stiglitz JE (2014) The price of inequality: How today’s divided society endangers our future. <http://www.pas.va/content/dam/accademia/pdf/es41/es41-stiglitz.pdf>.
- Tarafdar M, D’Arcy J, Turel O, Gupta A (2015) The dark side of information technology. *Sloan Management Rev.* 56(2):61–70.
- The Economist* (2016) Frankenstein’s paperclips. 419(8995):13–15.
- Thornhill J (2015) Brave new era in technology needs new ethics. *Financial Times*, <https://www.ft.com/content/dd328bf4-a25e-11e5-8d70-42b68cfae6e4>.
- Valett JD, McGarry FE (1989) A summary of software measurement experiences in the software engineering laboratory. *J. Systems Software* 9(2):137–148.
- Wei KK, Teo HH, Chan HC, Tan BC (2011) Conceptualizing and testing a social cognitive model of the digital divide. *Inform. Systems Res.* 22(1):170–187.
- Wolff EN (2010) Recent trends in household wealth in the United States: Rising debt and the middle-class squeeze—An Update to 2007. Levy Economics Institute Working Papers Series 159, New York University, New York. <https://ssrn.com/abstract=1585409>.
- Zuboff S (1994) *In the Age of the Smart Machine: The Future of Work and Power* (Basic Books, New York).
- Zuckerberg M (2010) How we’re fixing Facebook privacy. *Business Insider* (May 26). <http://www.businessinsider.com/how-were-fixing-facebook-privacy-2010-5>.