Sometimes Three Rights Really Do Make a Wrong: Measuring Cybersecurity and Simpson's Paradox[1]

Eric Jardine[2]

Assistant Professor, Political Science, Virginia Tech, and CIGI Senior Fellow

Paper presented to the 16th Annual Workshop on the Economics of Information Security, La Jolla, CA, 2017

Draft (June 30th, 2017).

Abstract: (223) Many of the over-time trends in available cybersecurity indicators appear to show that things are getting decidedly worse online. Yet many of these worsening trajectories might actually be based upon underlying trends that paradoxically show the situation in online security to be improving year over year. This peculiar reversal of fortune is known as 'Simpson's Paradox.' In this paper, I show that Simpson's Paradox emerges in the data on cybersecurity trends when three conditions exist: 1) the aggregate numbers are based upon data from definable subgroups; 2) the subgroups have differential propensities towards being hacked; and 3) the rate of expansion of these subgroups over time is mirrored by their vulnerability, with the most susceptible groups expanding the fastest. The near exponential growth of the IoT almost guarantees that these conditions exist online today. I test for Simpson's Paradox using a Monte Carlo simulation involving 2,002 iterations of randomly simulated data across both a Gaussian and a type-1 power law distribution. The tests show that a variant of Simpson's Paradox could easily be clouding our view of cybersecurity. The high possibility of a Simpson's Paradox in cyberspace entails that the worsening state of online security might not be as bad as many people think and that radically disjunctive policy reforms to 'fix' perceived cybersecurity challenges might not be warranted by the available evidence.

[1] I would like to thank Jason Kelly for his helpful comments through the various iterations of this paper. I would also thank the participants of the Inaugural Ostrom Workshop Colloquium on Cybersecurity and Internet Governance for their helpful insights. Mistakes remain my own.

[2] Author correspondence details: ejardine@vt.edu

In 1983, the Reagan Administration formed a National Commission on Excellence in Education to assess the state of K-12 schooling in America. The final report of the Commission, *A Nation at Risk*, charged that education was threatened by a "rising tide of mediocrity." The claim was supported, as the Commission saw it, by a steady decline in students scoring over 650 on the SATs, with an average drop of 50 points in verbal scores and a 40 point drop in mathematics.[1]

Based upon the findings of the report, both the Reagan and subsequent administrations radically overhauled K-12 education. The report led to changes in funding, curriculum, administration, classroom practices and to more standardized metrics of teacher and student evaluation. Many of these reforms were at once drastic and of questionable effect in terms of their ability to properly prepare students to face the rigors of an ever-changing economy. But, with test scores spiraling the drain, it seemed as though desperate times called for desperate measures.

The massive overhaul of educational policy and practice in America certainly seemed to be justified by the worsening state of student SAT scores as evidenced by the aggregated test outcomes. The problem was, student SAT performance was not actually in decline. Upon a second look at the numbers by researchers at Sandia National Laboratories, it turned out the observed nationwide decline in average SAT scores was masking some contrary underlying trends. When the numbers were disaggregated and observed at discrete categories of performance for students at the top, middle, and bottom of their high school classes, the trend in SAT scores actually showed modest improvement from 1975-1990. As a result, the Sandia researchers concluded, "the logical explanation for a decline in the combined average score [was] that the demographics of the students taking the exam [had] changed," with relatively poor-performing students making up a larger proportion of the sample over time.[2] Unfortunately, few had time for subsequent number crunching. At the level of official policy at least, the initial results stood, even though the statistical and evidentiary base of the massive (and often deleterious) reforms undertaken in the educational sector were based upon a faulty reading of the statistics.

It is a cautionary tale, one that cybersecurity researchers, practitioners and policymakers need to understand. On most accounts, the state of cybersecurity is also showing a worsening trend over time.[3] Each passing year seems to bring ever more breached records, hacked devices and new software vulnerabilities. Distributed denial of service attacks (DDoS) are growing in severity and even national power grids have become the tantalizing target of cyber intrusions, as was the case in Ukraine in December of 2015. These troubling trends have led former Homeland Security Secretary Tom Ridge to posit, in the fall of 2016, that "Notwithstanding the pain and horror associated with a physical attack. The potential for physical, human, and psychic impact with a cyberattack, I think, is far more serious."[4] Looking at these aggregate trends alone, the deteriorating state of cyber security seems to suggest the need for major new policy implementation and reforms, likely encompassing governments, private sector actors, insurance companies and even individuals.

Yet as was the case with educational performance in the US, the aggregate trends in cybersecurity might be misleading. If they are, then the perceived worsening in the state of online security could encourage governments, businesses and individuals to act in radically

disjunctive ways that are not in fact warranted by the evidence. In this paper, I use simulated data to show how simple it would be for the perceived overall worsening of cybersecurity to result from a version of what is known as Simpson's Paradox.[5] When the Yule-Simpson effect, as it is also known, is at work, aggregate trends can show a worsening situation, but the performance of distinct subgroups all demonstrate a trend towards an improving state of online security.

My core argument is a straightforward one. A Simpson's Paradox can manifest in cybersecurity measures if three highly plausible conditions are met. First, there are different groupings of hackable points online. Second, the groupings have different propensities towards being hacked. And third, the proportionate size of the groupings varies over time, with the number of worst performers growing fastest overall. As I show in more detail below, each of these three conditions quite plausibly obtains in today's online environment. Some people are savvy users who are relatively unlikely to be hacked. Others are more naïve users who are comparatively prone to being breached. And, increasingly, some hackable points are no longer even users, but are composed of the rank and file of notoriously insecure Internet of Things (IoT) devices.

Moreover, using a series of simulations each containing both a single graphical iteration and an additional 1,000 iteration Monte Carlo run, I show that Simpson's Paradox readily emerges when these conditions obtain. The results are also robust to diametrically opposed views about what the online world within which live is like. One set of simulations assumes the world of cybersecurity is normally distributed according to a Gaussian curve. The other assumes that the networks of the Internet form power law distributions. In Nassim Nicholas Taleb's language, the difference here is between a world of "Mediocristan" (Gaussian) and a the realm of "Extremistan" (power law).[6] Yet, in both potential worlds, the results support the idea that aggregate trends showing an increasingly insecure online environment could be the result of Simpson's Paradox rather than an actual overall worsening of cybersecurity. As a result, severely disjunctive policy responses aimed at improving cybersecurity might not be warranted, and may, as was the case with some aspects of early K-12 educational reform, actually make things worse.

By way of a roadmap for what follows, I first present a look at some aggregate-level trends in cybersecurity. The second section fleshes out in more detail the workings of Simpson's Paradox. The third section shows how the necessary assumptions for such a paradox to emerge in the realm of cybersecurity are easily met. The fourth section describes the simulated data for a Gaussian world and presents the results of both a single iteration graphical representation of the state of cybersecurity and a 1,000 iteration Monte Carlo simulation. Section five likewise describes the simulated data for a power law world and presents both a graphical and a 1,000 iteration simulation to test for the plausibility of Simpson's Paradox. The final section concludes by providing a discussion of the relevance of this finding for policymakers in the cybersecurity space.

Some Aggregate Trends in Cybersecurity

Looking at the raw numbers on cybersecurity, it is readily apparent why so many people have developed so much anxiety over the state of online security.

Privacy Rights Clearinghouse (PRC), for example, has collected a record of all disclosed data breaches in the US from 2005 to 2016. Like all data, the over-time trend is a bit noisy. As shown in Figure 1, the count of disclosed breaches, in particular, bounces around a fair bit, but nevertheless tends to show a trajectory towards ever more breaches over time. The average number of breaches in the first six years of the sample, for instance, is 381 breaches per year. The average for the second half of the sample, in contrast, is 498 disclosed breaches, making up a marked difference.

The trend in the number of breached records shows an even starker trajectory. From first to last, the number of breached records has swelled by 4,320 percentage points, increasing from 'just' 54,831,241 breached records in 2005 to 2,423,733,400 in 2016.[7] Mega breaches, such as the compromise of one billion Yahoo! account holders (counted in 2016 according to PRC data), have helped to grow these troubling numbers.

At the same time, a part of the change is also due to a worsening in the number of records compromised per disclosed data breach. In the first half of the sample, the number of stolen records per data breach was 403,889. However, by the last six years of the sample, the breach rate had grown to almost 2 million compromised records per data breach (1,900,666), a jump of slightly over 4.5 times.

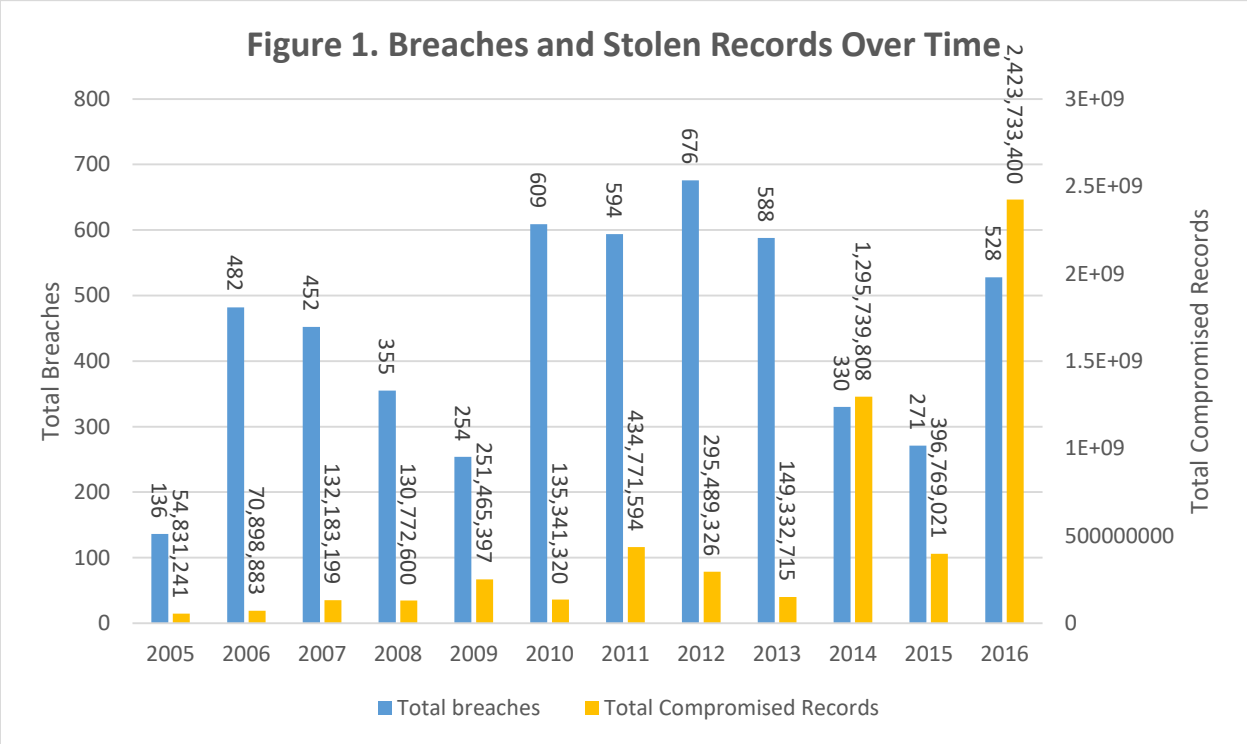

The simple average trends in Privacy Rights Clearinghouse data paint a picture that plausibly points towards the idea that cybersecurity is getting decidedly worse.[8] A number of factors could be at play, including changing data breach disclosure laws that require more companies, governments and non-governmental organizations to make public that they have suffered from a security incident. But even heuristically taking shifting regulatory regimes into account, the

numbers still seem to suggest a worrisome move towards more and more compromised personal details every year.

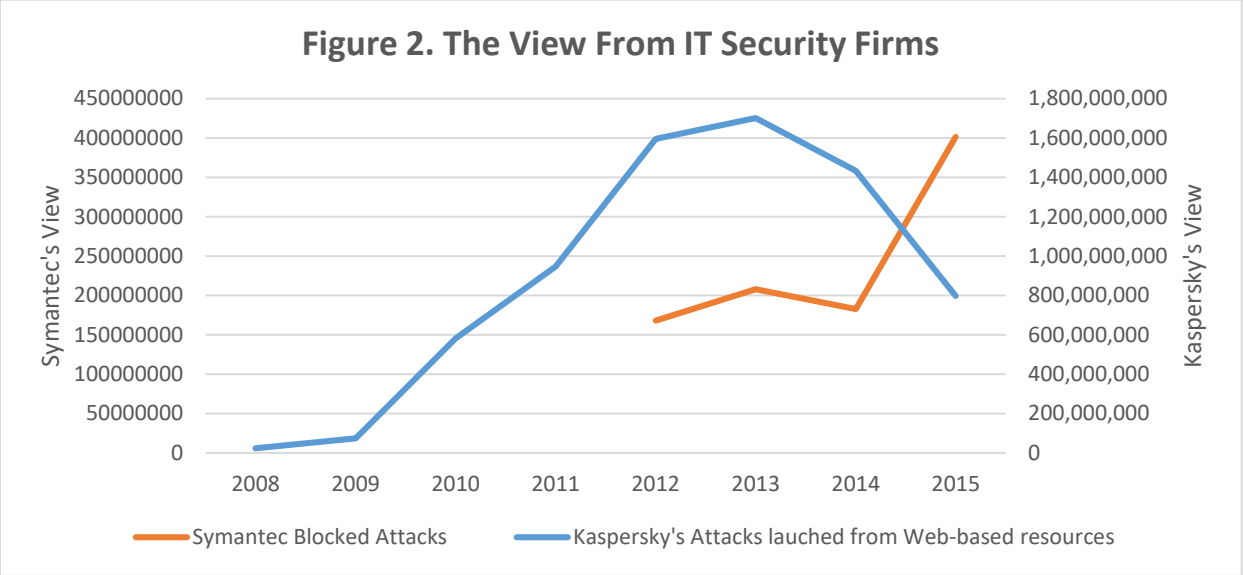

Other data sources tend to demonstrate a similar trend. IT security firms, who are in the business of protecting users while also drumming up new business, commonly issue annual reports documenting the poor state of cybersecurity. Kaspersky Labs, for example, issues an annual *Security Bulletin*.[9] In the bowels of the various reports over the years is a record of the observed number of attacks launched from online resources. Again, the trajectory, as depicted in Figure 2, shows a clearly worsening trends between 2008 and 2015, although attacks have declined since 2013. Indeed, even with 2 years of improving numbers taken into account, from the start (2008) to the end of the sample period (2015), the number of attacks has grown by 3,270 percent.

The antivirus and IT security firm Symantec observes a similar trend. Each year, they release an annual report known as the *Internet Security Threat Report*.[10] In the pages of these documents, Symantec provides a count of the number of cyberattacks that it has blocked for its customers over time, starting in 2012. With minor fluctuations aside, the trend is unambiguously positive: blocked attacks are going up. In this case, the count of blocked attacks has increased by 139 percentage points, rising from 167,900,000 blocked intrusions in 2012 to 401,500,000 in 2015.

There are, of course, a number of reasons to suspect that these trends in observed attacks might misrepresent the true extent of the problem. The correct measurement of cybersecurity indicators is a particularly perilous enterprise. Statistics based upon a survey of users, for example, can easily misrepresent the problem if the numbers follow a power law distribution rather than a normal Gaussian curve.[11] Statistics can also be widely biased by problems of underreporting.[12] Among firms, a lack of common metrics, both within a firm over time and across firms in a given sector, can cause serious problems in aggregate data.[13] The data on cybersecurity can also simply miss key facets of the problem, biasing estimates of cost, occurrence or severity.[14] Publicly available IT security vendor data, finally, can be particularly error prone, as these firms

are in the uncomfortable position of doing research on a phenomenon the occurrence of which they depend upon for their livelihood.[15]

Additionally, even if the count of costs, attacks, vulnerabilities or what have you is correctly collected and measured, a failure to normalize these statistics around the growing size of the Internet ecosystem can still lead people to mistakenly conclude that things are far worse in the realm of online security than they actually are.[16] Increases in the total number of attacks as documented in the data above might not necessarily mean, therefore, that things are actually growing worse overall.

Yet even when the number of attacks are normalized in appropriate ways, the overall aggregate trends still often shows a worsening trajectory.[17] The perception (and practice, in some cases) of worsening numbers has led to a broad-based deterioration in the sentiment of cybersecurity practitioners. The Index of Cybersecurity, for instance, has documented a steadily worsening perception of systemic risk among cybersecurity professionals, which has increased from under 1,500 in May 2012 to 3,640 as of April 17, 2017.[18] Across a variety of metrics, then, the aggregate situation in cyberspace appears to be pretty grim. Data breaches are becoming more common and affecting more people, all while various forms of online attacks continue to increase, potentially leading to millions of people being infected with malware of various stripes each and every year.

Based upon an observation of just these aggregate trends, it would be fair to conclude that the deteriorating state of cybersecurity clearly warrants significant, disjunctive and far-reaching policy reform. By simply observing that the average trend is poor, it is easy to conclude that we need to reach deep and reform everything we can in order to make our digital ecosystem more secure. But, as was the case with educational reform following the *Nation at Risk* report, these radical policy reforms might not be warranted. The aggregate trends might be lying to us.

Simpson's Paradox

Despite being something many people have never even heard of, Simpson's Paradox—otherwise known as Yule-Simpson effect or aggregation bias—can very easily "trap the unwary."[19] For something so far from the public imagination, it continually manifests in situations as diverse as university graduate admissions, mortality rates among smokers and non-smokers and price elasticity in the economy—and, of course, educational outcomes.[20] When it does emerge, as we saw with the example of the *Nation at Risk* report, aggregate negative trends can actually be based upon underlying positive outcomes, and, conversely, positive trends at the aggregate level can actually be based upon underlying negative ones.

A medical study assessing the effectiveness of various kidney stone removal techniques is one famous example of the logic of Simpson's Paradox in action.[21] The study's aim was to determine which type of surgery was the most effective at eliminating kidney stones. To isolate for the effectiveness of each treatment, the researchers broke the study population of 1,052 patients with renal calculi into two groups: those with kidney stones that were greater than 2 cm and those

whose stones were less than 2 cm. Had they not, a significant portion of their results would have fallen prey to Simpson's Paradox.[22]

While a non-invasive procedure known as extracorporeal shockwave lithotripsy (ESWL) was the most effective treatment overall, the peculiarities of the Yule-Simpson effect emerge quite distinctly when comparing the aggregate and disaggregate numbers for open surgery techniques and percutaneous nephrolithotomy.

Table 1. Aggregate Success Rates of Open Surgery and Percutaneous Nephrolithotomy

	Total N	Number of Successful Cases	Success Rate	Best Treatment
Open Surgery	350	273	78%	No
Percutaneous Nephrolithotomy	350	289	83%	Yes

Table 1 presents the aggregate results of a comparison of these two treatment methods. Open surgery was effective in 273 instances, representing a success rate of 78 percent. Percutaneous nephrolithotomy, in contrast, was successful in 83 percent of the treatments. These results come out quite clearly in favor of percutaneous nephrolithotomy as the more effective treatment method. But, these aggregate trends—as we saw before in the case of educational outcomes—mask some confounding underlying dynamics.

Table 2. Disaggregate Success Rates of Open Surgery and Percutaneous Nephrolithotomy

		Total N	Number of Successful Cases	Success Rate	Best Treatment
Group 1 (<2cm)	Open Surgery	87	81	93%	Yes
	Percutaneous Nephrolithotomy	270	234	87%	No
Group 2 (>=2cm)	Open Surgery	263	192	73%	Yes
	Percutaneous Nephrolithotomy	80	55	69%	No

Table 2 presents the disaggregated results for the same comparison. When the cases are divided into the researcher's study groups (again, those with kidney stones less than 2cm (Group 1) and those with stones greater than or equal to 2 cm (Group 2)), the most effective treatment is inverted. In contrast to the aggregate results, open surgery emerges as the winner. For Group 1, for example, open surgery is successful 93 percent of the time, compared to an 87 percent success rate for percutaneous nephrolithotomy. For Group 2, open surgery is effective in 73 percent of the cases, compared to a success rate of just 69 percent for the alternative method. In short, open surgery, which was less successful in the aggregate, emerges as the more successful

method when the results are disaggregated and considered conditional upon the size of the patient's kidney stones.

The peculiar reversal emerges because the size of a person's kidney stones is a so-called "lurking confounder." A lurking confounder of this sort has to have two characteristics in order for a Simpson's Paradox to emerge. In the exemplar language of this study, it has to be correlated with both the type of surgery employed (X) and the outcome or effectiveness of the surgery (Y).[23] The size of a person's kidney stones easily fits the bill.

First of all, kidney stone size is clearly correlated with treatment options. Assignment to the various treatments in this case was not done at random, but rather availed upon the prevailing medical judgement of the doctors involved. This means that the treatment that a patient received was significantly related to the size of their kidney stones. For example, only 25 percent of the 350 instances of open surgery were used in Group 1 (small kidney stone) cases. The remaining 75 percent of open surgeries were used in the more severe cases of kidney stones over 2cm in diameter. Percutaneous nephrolithotomy displays the opposite trend. Here, the largest proportion (77 percent) of cases are clustered in group 1 containing those individuals with kidney stones under 2 cm in diameter.

Kidney stone size is also correlated with success rates regardless of employed treatment type. In this particular study, success was measured as being free from kidney stones three months after the treatment. Larger, more complex cases should be harder to treat, everything else being equal, and the numbers tend to bear this logic out. Of the 357 Group 1 cases involving both treatment types, some 315 (88 percent) were successful. In contrast, within the more severe cases in Group 2, only 247 of the 343 treatments were effective, marking a relatively poor success rate of only 72 percent.

In sum, when a lurking cofounder that is correlated with both the outcome and the independent variable of interest exists, any observed aggregate relationships become susceptible to Simpson's Paradox. Trends in one direction at the aggregate level can even reverse when the data is disaggregated into constituent categories. As I show in the next three sections, the perils of Simpson's Paradox could easily apply in today's world of cyberspace.

The Plausibility of the Three Necessary Conditions for Simpson's Paradox in Cybersecurity

As shown in the initial section, aggregate trends in cybersecurity tend to be pointing in a decidedly negative direction. US data breaches seem to be growing in both frequency and severity, while globally observed malicious attacks are mounting. With the caveats about the importance of proper measurement methods and normalization aside, it looks for all intents and purposes like efforts to enhance cybersecurity have been, so far at least, a largely losing battle. At the very least, this negative sentiment is what is largely believed by many policymakers in government and the private sector.[24] But what if the Yule-Simpson effect is at work in cybersecurity?

Simpson's Paradox readily emerges in cybersecurity if three conditions hold:

- First, there are definably distinct groups of hackable points online;
- Second, these hackable points have different propensities towards being hacked;
- And, finally, the growth of these groups over time is uneven, with the most susceptible category being the fastest growing segment.

Each of these conditions are exceedingly plausible given the Internet's rapid commercialization since the launch of the World Wide Web in the early 1990s, particularly with the recent turn towards more interconnected devices than humans. Simply put, Simpson's Paradox might be clouding our view of cybersecurity and things might be far better than they at first appear on the surface.

Let's begin by fleshing out the plausibility of each of the three conditions, before looking at how their combined effect results in a Yule-Simpson effect in cybersecurity. Take the first condition, that there are definably distinct groups of hackable points online. As we saw before in the kidney stone treatment example, these groups cannot be any old set of groupings. Depending upon how you slice the pie, there are already plenty of definable groups online, such as trolling subcultures,[25] hacktivist collectives such as Anonymous,[26] globe-spanning cybercriminal gangs,[27] and a myriad other formal and informal associations. But, the important groupings for the issue of measuring cybersecurity correctly have be potential victim categories. Victim categories, in other words, are confounders; that is, they have some relationship to both the likelihood that a person or device will be hacked and the growth of the Internet over time. Groups of this sort are easy enough to conceptualize. With enough information on the disposition and behavior of users (broadly defined), it would be possible to divide the world up into three potential target types or cybercrime victim categories: savvy users, naïve users and Internet of Things devices. The existence of such potential grouping satisfies the first condition.

Obviously, each of these points could be hacked, but their propensity or vulnerability toward external intrusion would significantly vary. This variation satisfies the second condition. Savvy users, for example, would be individuals with a certain amount of technical expertise, perceptions of self-efficacy or knowledge of IT security threats.[28] Such users would take at least some precautionary steps to protect themselves online, such as using antivirus software, avoiding certain websites, and knowing, in general terms at least, what sort of things to avoid doing in email.

Savvy users are likely the smaller subset of all human users. For instance, one survey of computer users, conducted by the Organization for Economic Co-Operation and Development (OECD), found that only 31 percent of people are tech savvy. To determine this figure, the survey had some 215,942 people across 33 countries complete 14 general computer-based tasks. Only 5 percent fell into the highly competent category at the top of the pile, while another 26 percent fell into the moderately skilled camp.[29] Plausibly, then, the sum of the two (31 percent of users) fall into the savvy user category at the end of 2015. The number of savvy users can be further extrapolated into the end of 2017. Since all users must be either savvy or naïve, the remaining 69 percent could be classified as more naïve users of networked technologies in 2015 and so forth.

The other definable grouping of hackable points online is not human at all, but is composed instead of Internet-enabled devices. The explosive growth of the IoT—which I mark as starting in 2008 when the number of Internet-connected devices surpassed the number of actual Internet users—provides a host of targetable points that malicious actors can exploit. The devices themselves are, moreover, notoriously vulnerable to external intrusion. One 2015 report by Symantec on *The Insecurity in the Internet of Things*, for example, found ten significant security issues with IoT home products. These issues included worrying results such as a severe lack of mutual authentication procedures between device clients and servers; limited use of password protection best practices such as two-factor authentication, the inability to change weak default passwords and the inadequate use of password lock-out measures to slow brute force attacks; as many as 10 security issues in the 15 portals used to remotely control IoT devices; and a litany of common web application vulnerabilities plaguing the various IoT cloud platforms.[30] Another study by HP conducted in 2014 likewise found that upwards of 70 percent of IoT devices were vulnerable to being hacked.[31] The rapid expansion of the Mirai malware that turns IoT devices into botnet zombies—causing a bunch of DVRs and babycams to take down Netflix in the fall of 2016—is a good example of the vulnerability of this segment of the overall ecosystem.[32] In short, it is fair to say that the lack of security built into IoT devices makes them even more vulnerable than most naïve human users employing laptops, desktops and mobile devices and a only a modest dose of common sense.

The last condition that needs to be present for Simpson's Paradox to emerge is that these three groups of distinctly hackable points must grow at different rates over time, with the more vulnerable camps expanding faster than more secure users. Publicly available World Bank data on Internet users and data on IoT expansion from Statista Research gives us a starting point for the numbers.[33] Subsequently applying a couple of simplifying assumptions makes it easy enough to plot out the relative growth of each grouping from the first quarter of 1990 to the end of 2017.

Users, first of all, need to be split between the savvy and naïve categories. Inevitably, the proportionate division between savvy users and naïve users is not static over time. It is quite plausible that early adopters of the Internet were almost by definition tech savvy people—which is not to say that they could not fall prey to problems, only that they had more understanding of the mechanics at play than later users. If we assume for the sake of argument that 99 percent of the Internet's original users in 1990 were savvy and that the 31 percent of people who could complete numerous computers tasks are the savvy proportion of the population in 2015, then we can employ a simple linear interpolation equation to show the proportion of savvy and naïve users per quarter between Q1 1990 and the end of 2015 and a linear extrapolation function to determine the proportion of savvy users at the end of 2017.[34] Obviously, parts of the Internet's expansion have been decidedly non-linear and interpolation over long time frames is inherently fraught with perils, but the general notion that one group expands as the other contracts is well served with the simple linear interpolation procedure.

Plotting the growth in the IoT from it 'birth' in 2008 to the end of 2017 is a bit simpler. Statista Research has estimates for the total annual number of IoT devices from 2012-2017. Assuming

that the IoT was born in Q1 2008 when the number of IoT devices was equal to the number of users plus one, then here again linear interpolation can fill in the quarterly gaps.

Figure 3 plots out the actual quarterly trends in the absolute number of hackable points in each grouping from Q1 1990 to the end of 2017. The central takeaway is that all three are clearly on the rise, but the rate of increase is decidedly different. The IoT is growing fastest, handily beating out both types of users. Among the human user category, the total number of savvy users is getting decidedly outpaced by the increase in naïve users, particularly after around the end of 2004. In more precise terms, from 2008 when the IoT entered the sample until the end of 2017, the IoT grew some 1631 percent, Naïve users grew by 310 percent and savvy users grew by only 34 percent. In short, the third necessary condition for the Yule-Simpson effect to be at play is plausibly satisfied.

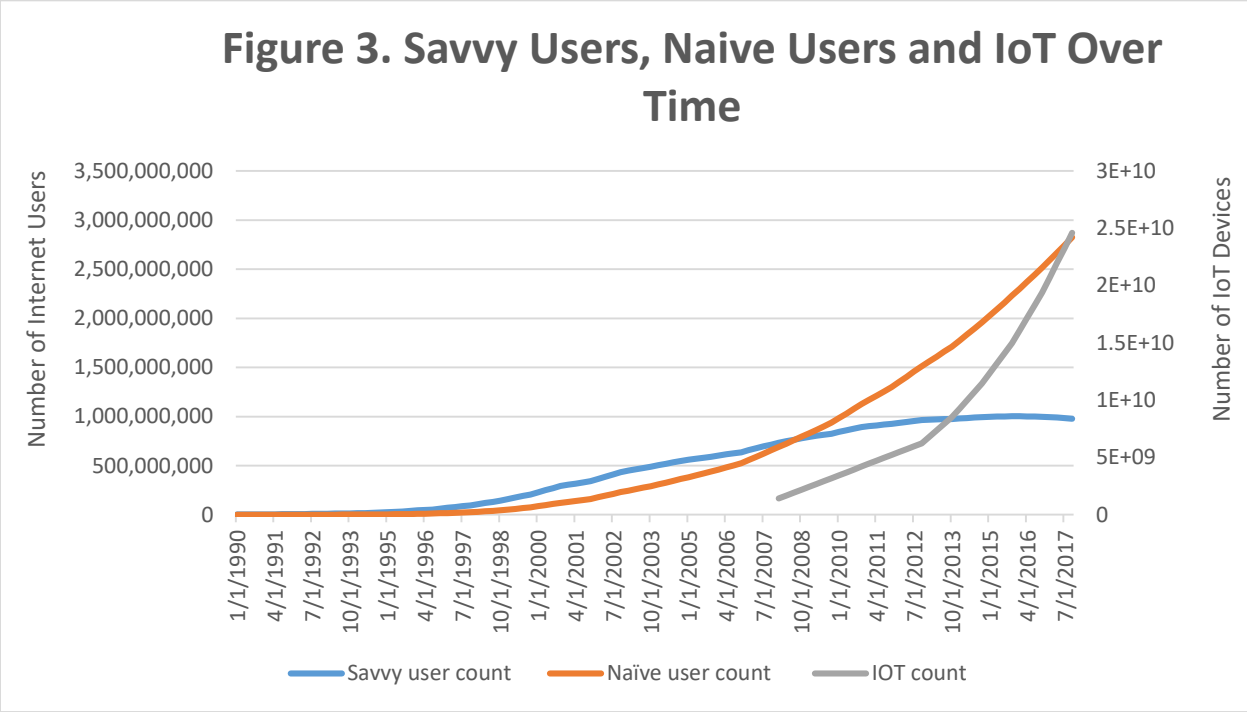

Given that the three conditions necessary in order for the Yule-Simpson effect to be distorting our view of cyberspace are plausibly met, the remaining task is to quantify some of the implicit notions above and employ simulated data in order to look to see if a deteriorating overall trend can actually mask three distinctly improving trajectories. The following discussion in sections four and five does this in two ways and across two different views of reality (a normal distribution and a power law distribution). First, for each scenario, I present a single iteration of the simulated data and plot the results graphically to highlight the point. Second, I use a Monte Carlo simulation to determine the results of the simulations over 1,000 iterations of the data. In both cases, the results support the idea that a Yule-Simpson effect could easily be at work in cybersecurity.

Cybersecurity in a 'Normal' World

This section details how when the three conditions hold, Simpson's Paradox can cloud our view of cybersecurity in a world that is characterized by normal, Gaussian distributions. To do so, I employed Excel's (pseudo)random number generator to simulate a normally distributed number expressing the proportion of each category of user will be hacked in a given quarter. This number can be effectively understood at a conceptual level as a proportion of users breached in each category in each quarter.

Table 3. Summary Statistics for The Graphically Presented Simulated Data

	Improvement Rate	Min	Max	Mean	Median	Standard Deviation
Savvy Users	0.1	0.29	0.50	0.40	0.40	0.044
Naïve Users	0.1	0.44	0.62	0.54	0.53	0.042
IoT	0.1	0.59	0.73	0.66	0.66	0.032

Table 3 presents summary statistics for the data used in the single iteration of the simulations. The min, max, mean, median and standard deviation categories are pretty self-explanatory and capture the basic statistical dimensions of the simulated data. To generate data that ensured differing hack proportions for each category of user or device, the mean in the random generator function was set at 45 for savvy users, 60 for naïve users and 70 for IoT devices. The standard deviation in the random number generator was 3.33 for each, creating a set of random numbers that are effectively bounded in the range of 35-55 for savvy users, 50-70 for naïve users and 65-85 for IoT devices. The overlap in these ranges captures the invariably fuzz boundaries that would mark membership in each category.

The rate of improvement column captures the notion that each group is getting modestly better at staying safe online over time. The improvement rate is determined probabilistically. Each group has some probability (0.1) that the randomly generated hack proportion value for each quarter will be discounted by some additional count value. The discount rate for each group remains constant over the range of the data, but the size of the associated count value increases in size by integer values starting at 0 and ending in the last quarter of 2017 at 112.

With this sense of the dimensions of the single iteration of the simulated data in a Gaussian world under our belts, we can turn to the plotted data from the single iteration to observe whether Simpson's Paradox could be clouding our understanding of cybersecurity.

In line with the discussion above about the aggregate cybersecurity trends as observed by Privacy Rights Clearinghouse and various IT security firms, an initial one-off run of the simulated data in aggregate shows a decidedly worsening situation over time. The aggregate simulated data are based upon the weighted average of the proportion of users hacked in each quarter for the various subsamples for each quarter for each year. For example, in 1990, when we assumed savvy users made up 99 percent of the user base of the Internet and when the IoT was

not yet in existence, the aggregate value for Q1 was heavily weighted toward the hack proportion of savvy users (savvy users hack proportion = 0.44, the naïve users hack proportion was 0.51 and aggregate total hack proportion = 0.437). As more naïve users flooded onto the scene, the aggregate hack proportion shifted each quarter to reflect the weighted average of users in the two categories, as it continued to do after 2008 when the IoT became a source of hackable points in the Internet ecosystem.

Figure 4 plots the aggregate (weighted average) relationship between the growth of the Internet over time (x) and the aggregate estimated proportion of users that are hacked each period (y). The plotted quarters are overlaid with a linear trend line that is meant to be indicative of the directionality of the movement of the data over time and not a more analytical regression. In the early 1990s, the average hack proportion is around 42 percent. By the end of the simulated data in 2017, the average is well over 63 percent, indicating a decidedly worsening picture of online security. Based upon this simulated data, as with the real trends in data breaches and observed attacks noted above, the situation in cyberspace appears to be getting worse over time, at least according to the aggregated numbers.

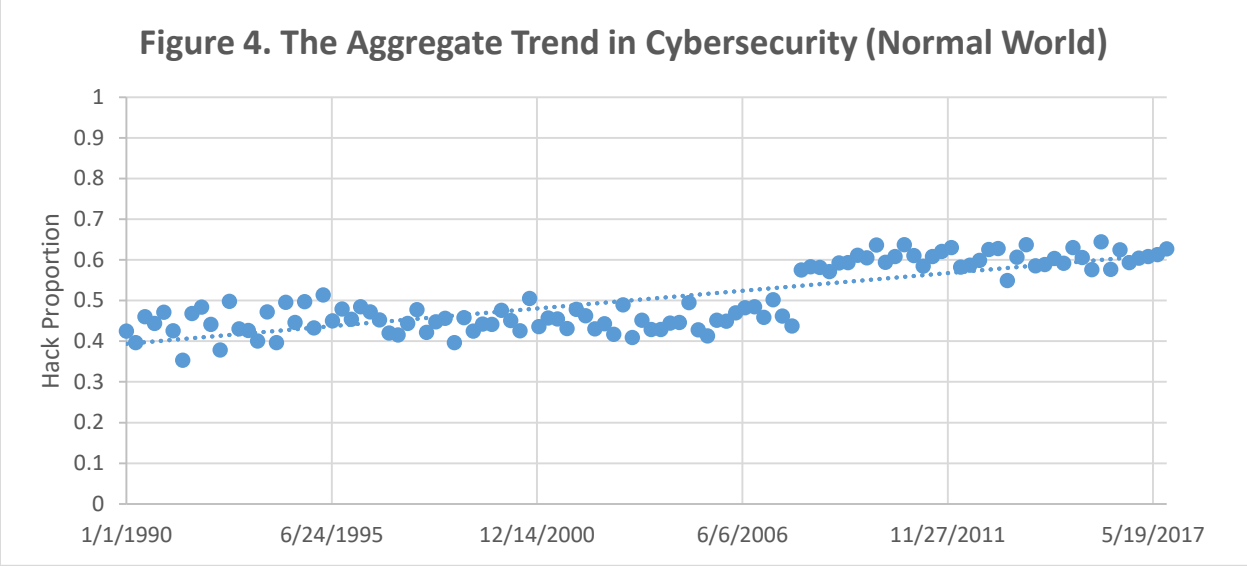

However, once the aggregate data is decomposed into its constituent groups, the seemingly worsening situation reverses itself. Figure 5, for example, plots the relationship between the expansion of the Internet over time and the proportion of savvy or sophisticated users that are hacked in each quarter, again overlaid with a linear trend line. Unlike the trend in the aggregate data, the trend for savvy users is decidedly negative. In 1990, the average hack proportion for savvy users was just over 42 percent. By the end of the simulated range in 2017, the average hack proportion had declined to closer to 29 percent.

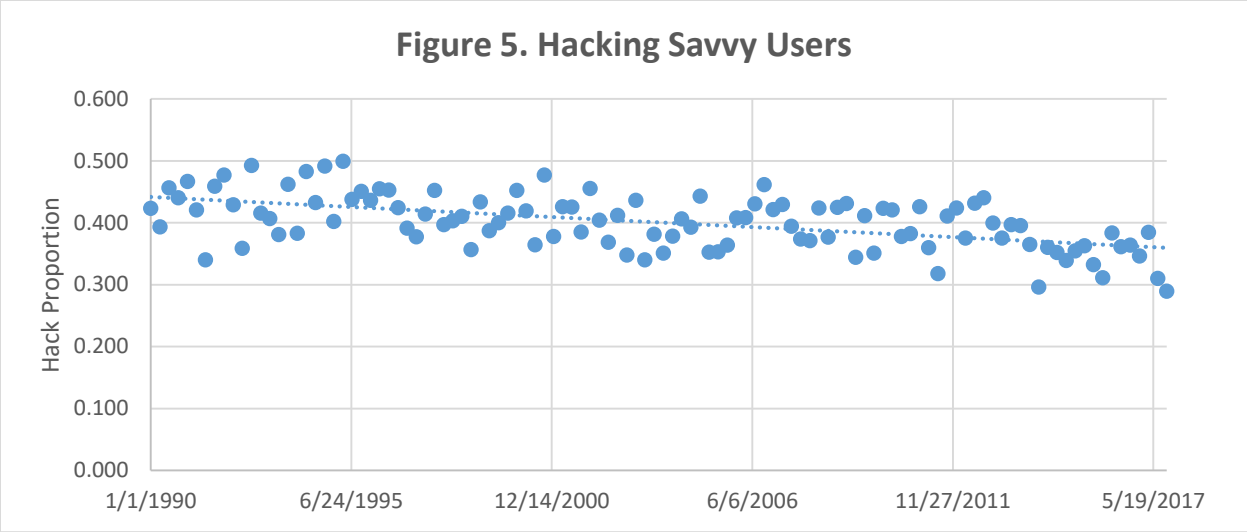

A similarly negative trend is found in the naïve user subcategory. Figure 6 plots the simulated data. Like in the savvy user subsample, the trend between the expansion of the Internet over time and the proportion of users that are hacked in each quarter is negative. In 1990, the proportion of naïve users that are hacked according to the simulated data was roughly 61 percent. By 2017, the proportion of hacked users declined by over 16 percentage points, falling to 45 percent of the category population. According to these data, naïve users—as with savvy users—are actually becoming safer online over time, despite what the aggregate trend indicates.

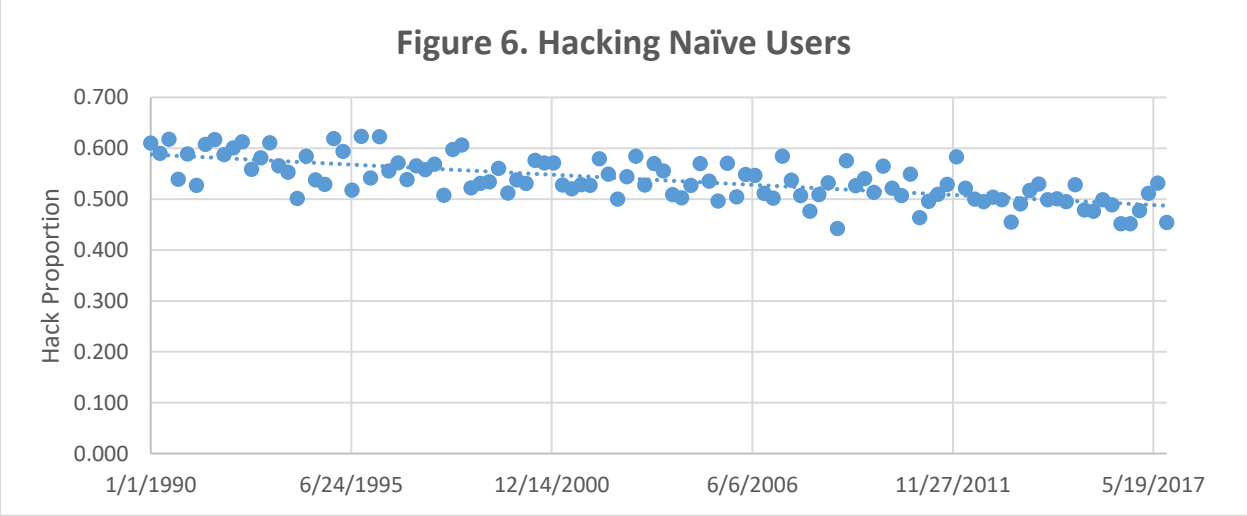

Figure 7, finally, plots the IoT subsample. As in the other two instances, the relationship between network expansion over time (x) and the simulated proportion of IoT devices that are hacked (y) is negative. Starting in 2008, when the IoT was brand new and (even more) riddled with insecure code, around 73 percent of devices were hacked in this one-off run of the simulated data. By the end of the sample in 2017, the proportion of hacked IoT devices declined to slightly less than 66 percent, which would accommodate for people gradually coming to grips with the need for security by design in IoT devices. The decline for the IoT is less pronounced than the other categories due to the shorter time frame for its existence. Had the IoT existed since 1990 as well,

it too would has seen a much sharper decline in the proportion of devices that are hacked each quarter.

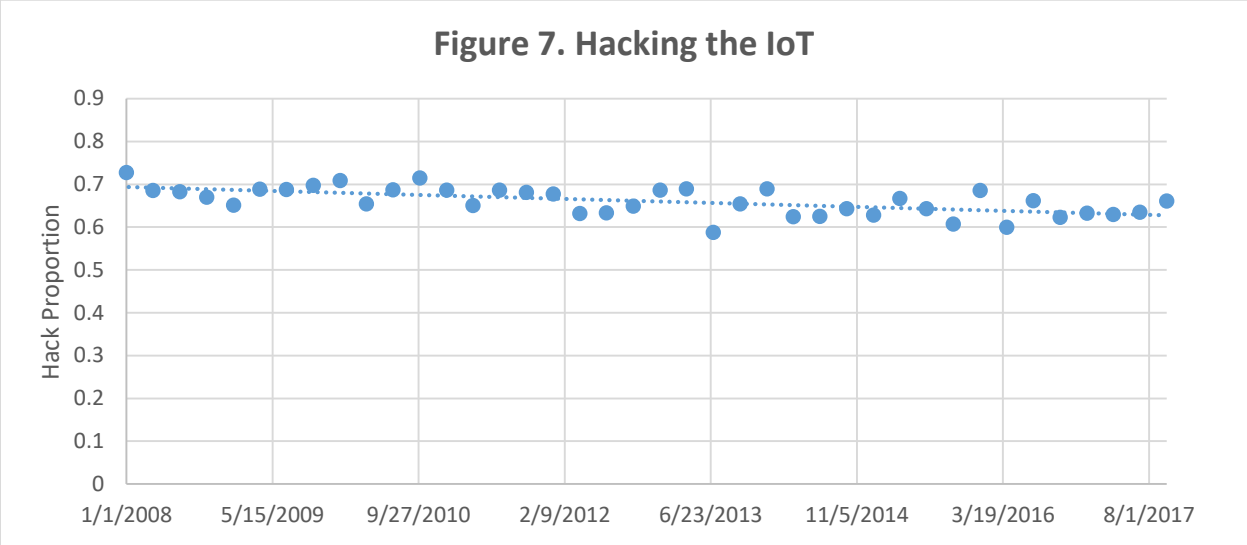

The figures capture the output of a single iteration of the simulated data in a normally distributed world. The basic implication is that a variant of the Yule-Simpson effect could indeed be increasing the public's perception of growing threats to their security in cyberspace. However, randomly generated data over a single iteration can show almost anything and could easily be quite far from the norm. To correct for the potential randomness and extremity of a single iteration of the data, I ran a Monte Carlo simulation to produce an average estimated hack proportion for each category of hackable points over 1,000 iterations of the data.

The results of the Monte Carlo simulation are in line with those of the single iteration presented in the graphs above. Table 4 presents the average hack proportions for the 1,000 iterations of the data. For ease of presentation, the 112 quarterly averages for each category of hackable point over the 27-year range are broken into quartiles. Each quartile encompasses seven years or 28 quarterly periods. The last column in the table presents the percentage point change in the proportion of users from each group that are hacked from the first to the last quartile.

As detailed in Table 4, all three of the user categories experience a decline over this period over the 1,000 iterations of the data. Savvy users and naïve users, both of which are in the sample from the very beginning, decline the most, falling 9 percentage points each.[3] The IoT category, likewise, falls by a relatively modest 2 percentage points. In contrast to these declining numbers, the average proportion of units that are hacked over the 1,000 iterations of the data for the aggregate trend is positive, increasing over the study range by 16 percentage points. Like the results of the single iteration of the data presented above, negative trends aggregate into an overall positive trajectory.

[3] Since the decline rate is the same across each user category, having both savvy users and naïve users decline in the same amount is to be expected. If either the rate of decline or the number of years that one of the categories was in the sample was to change, the gap between the two categories would grow more pronounced.

Table 4. Simpson's Paradox Over 1,000 Iterations of the Data (a Gaussian World)					
	Hack Proportion (1990-1996)	Hack Proportion (1997-2003)	Hack Proportion (2004-2010)	Hack Proportion (2011-2017)	Over Sample Percentage Point Change
Savvy Users	44%	41%	38%	35%	-9%
Naïve Users	59%	56%	53%	50%	-9%
IOT	N/A	N/A	67%	65%	-2%
Aggregate Trend	45%	45%	52%	61%	16%

In short, the simulated data presented here suggest that common observations of the worsening state of cybersecurity could be falling prey to Simpson's Paradox. In aggregate, the simulated results show things getting worse, with a higher proportion of people and devices getting hacked over time (potentially up to 16 percentage points more likely, according to the Monte Carol simulations). The trends in the underlying data, however, are decidedly different. For each discrete subgroup (savvy users, naïve users and IoT devices), the proportion of users and devices that are hacked each quarter is going down. Paradoxically, these results imply that cyberspace is getting safer, even as it seems at a more general level to be getting more insecure. Simpson's Paradox may well be striking again.

Simpson's Paradox in a Power Law World

The discussion above assumed that the world of cybersecurity is neat and tidy, characterized mostly by a series of normally distributed curves. The implication of this assumption are manifold. Foremost among them is the idea that variance in the distribution of cybersecurity incidents is predictably constrained. Because of the underlying math that characterizes Gaussian distributions, 68 percent of observations in the simulated data above fall within one standard deviation of the mean, 95 percent of observations within two standard deviations and, of course, 99 percent fall within three standard deviations. This regularity of occurrence is possible in some realms (human height being a prime example), but the networks of the Internet tend to form into so-called power law distributions, suggesting that cybersecurity might not be normal in the statistical sense of the term.[35] Indeed, extrapolation from cybersecurity survey data is confounded by the underlying power law distribution of cybercrime.[36]

When power laws govern the day, distributions develop large tails and are highly prone to outliers, which can easily be so huge that they make summary statistics such as the mean a lot less precise.[37] A number of natural phenomenon form power law distributions, ranging from the use of words in the English language to earthquakes and academic article citations. Power Laws, also known as Pareto Distributions, can be so skewed that they are often loosely governed by the so-called 80/20 rule, where some 80 percent of the variance in the outcome of concern can be driven by just 20 percent of the observations. In short, a world that is governed by power law distributions is far more prone to extremes than the world of nice, bell-shaped Gaussian curves.[38]

To determine whether Simpson's Paradox can emerge in a world of extremes, I simulated a type-1 power law distribution again using Excel's (pseudo)random number generator.[4] In this simulation, I used some proportion of each category of hackable points as the Xmin value. I then divided the outputted number from the full equation by the total population for each respective category in each quarter. This essentially produces a number that can again be read as a proportion of users or devices breached in each quarter. At times, the proportion exceeds one, indicating that, effectively, all users and then some are hacked in a given quarter. A proportional value of 5, for instance, means that every user in the category has been hacked five times in that quarter, assuming that hacks are uniformly distributed within a category. As we shall see, a power law world is much more chaotic and prone to variance than the Gaussian world. But, when the three conditions hold, Simpson's Paradox still emerges to once again cloud our view of cybersecurity.

Table 5. Summary Statistics for Cybersecurity in a Power Law World

	Xmin	α	Min	Max	Mean	Median	Standard Deviation
Savvy Users	Savvy Users * 0.4	3	0.33	1.02	0.50	0.45	0.15
Naïve Users	Naïve Users * 0.5	2.5	0.39	5.25	0.77	0.57	0.57
IoT Device	Devices * 0.6	2	0.45	6.60	0.98	0.67	1.03

Table 5 presents some preliminary summary statistics for the single iteration data used below. The Xmin column shows the minimum number of users breached in each period. In order to fulfil the assumption that the groups have differential likelihoods of being hacked, the proportion of each user or device category that can be hacked increases from 0.4 in the savvy user case to 0.6 with IoT devices. The α column shows the value on the exponent. In this case, the exponents decrease slightly across the categories of users. As the exponent shrinks, the variance in the distribution tends to increase, so an α of 3 for savvy users is less volatile and prone to extremes than an α of 2 for IoT devices.

Power law distributions are marked by extremity by their very nature. This feature of the distribution type is well represented in the gap between the minimum and maximum values for all categories of users and devices, as shown in Table 5. The min for savvy users, for instance, is 0.33, suggesting that a minimum of 33 percent of savvy users are hacked in a given quarter. The maximum is 1.02, which suggests that every user and then some (2 percent) could be hacked.

[4] Since Excel does not have a built in function to generate power law curves, I use the following form: (xmin * (RAND()*(1+β/112))^(-1/α))/p, where xmin is the minimum value of x (number of users in each category who are hacked per quarter), RAND() generates a uniformly distributed variable between 0 and 1, β is a count variable ranging from 0 in period t-1 to 112 in the last quarter 2017, α is the exponent value and p is the whole population in that category of hackable point in that period. This form of the equation essentially produces a number that can be read as the proportion of users in each category of user or devices that have been hacked in a given quarter. I am grateful to John Coleman for helping me figure out how to model this sort of world.

The most hackable category, IoT device, is even more extreme, with a minimum hack proportion of 0.45 and a maximum of that is rough 14.5 times larger at 6.60, suggesting that every devices is breached 6 plus times in one quarter. With maximum values like this one, it is clear that either some segments of the system can be breached repeatedly or every device is hit more than once (or a combination of both) in each quarter. This growing variance across the hackable groups is a result of the decreasing size of α.

The other feature of power law distributions that crops up in stark relief is the gap between the mean and the median. With a normal distribution, the two statistics represent effectively the same thing in practical terms (see Table 3 above). In a world shaped according to power law distributions, the two statistics reveal very different facets of the underlying numbers. The mean is artificially high (and hugely unstable) because it is being driven up by massive outliers that add a disproportionate amount to the summed value of the outcome at hand. The median—as medians always do, splits the sample so that 50 percent of the observations are above the value and 50 percent are below the stated value. In some ways, the gap between the two statistics really shows just how much of the outcome is being driven by extremes. For example, there is a 30 point spread between the mean and the median for the category of IoT devices. This suggests that, as again befits the role of a smaller exponent, the outcome of interest is highly prone to huge spikes. Regardless, despite the fundamentally different mechanics that are at work in the potential power law world of cybersecurity, Simpson's Paradox can still emerge to cloud our view.

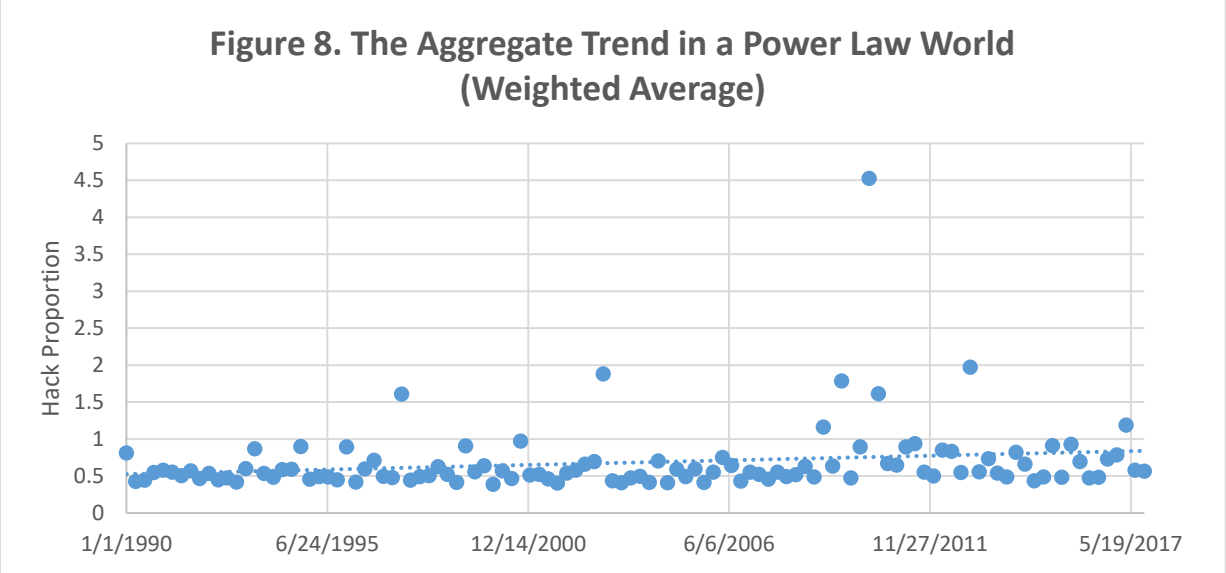

Figure 8, for instance, presents the results of a single iteration simulation of the aggregate trend in cybersecurity in a power law world of extreme values. As was the case with the aggregate trend in the data within a world governed by Gaussian distributions, the aggregate trend in the proportion of users and devices that hacked is again positive, increasingly from roughly 55.9 percent in 1990 to 78.2 percent in 2017. The variance within the data is clearly far more pronounced than what happened within the Gaussian world. While the proportion of users or devices that are hacked in the world of normal distributions tends to cluster fairly tightly around

the trend line, the numbers here swing wildly, sometimes hitting a value that is at least four times as much as the fitted linear trend line.

Disaggregating the data reveals both the continuance of the high variance world of power laws and a series of improving trends in the proportion of people and devices that are hacked in each quarter. Figure 9, for example, plots the disaggregated data for the savvy user category. Despite the worsening aggregate trend, the general trajectory of hacking amongst the category of savvy users is decidedly negative.

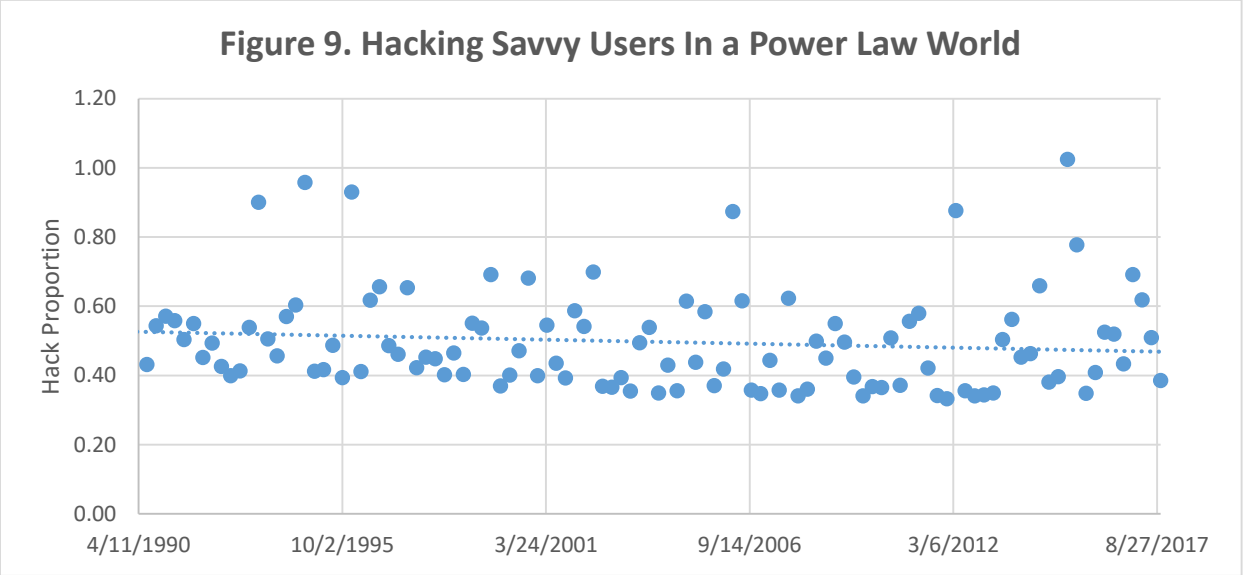

Figure 10 depicts the outcome of the single iteration of the data for the category of naïve users. Despite the high variance in outcomes, exemplified by the hugely positive values on specific quarters, the overall trend is again negative. In this case, the proportion of naïve users that are hacked in 1990 is around 83 percent, while the final proportion of hacked users for 2017 is essentially 73 percent, (although it is hugely important to stress that due to the nature of power laws, these averages are very prone to change and should not be read as having too much concrete meaning). Figure 11, finally, repeats the procedure with the simulated data for the category of IoT devices. The disaggregated trend is again negative, standing in stark contrast to the overall aggregate trend.

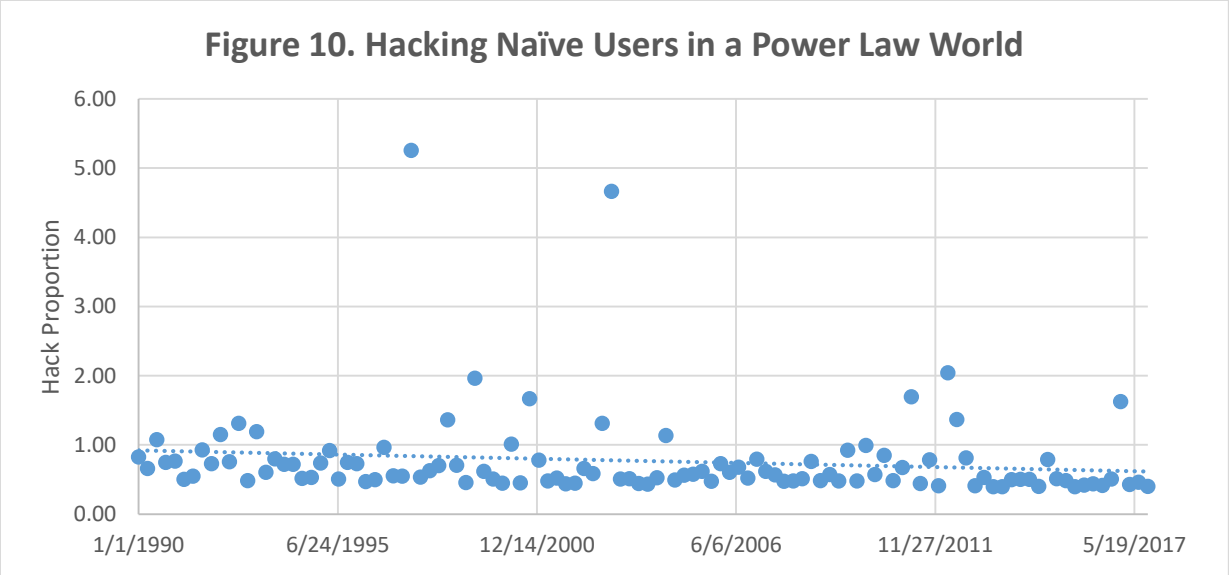

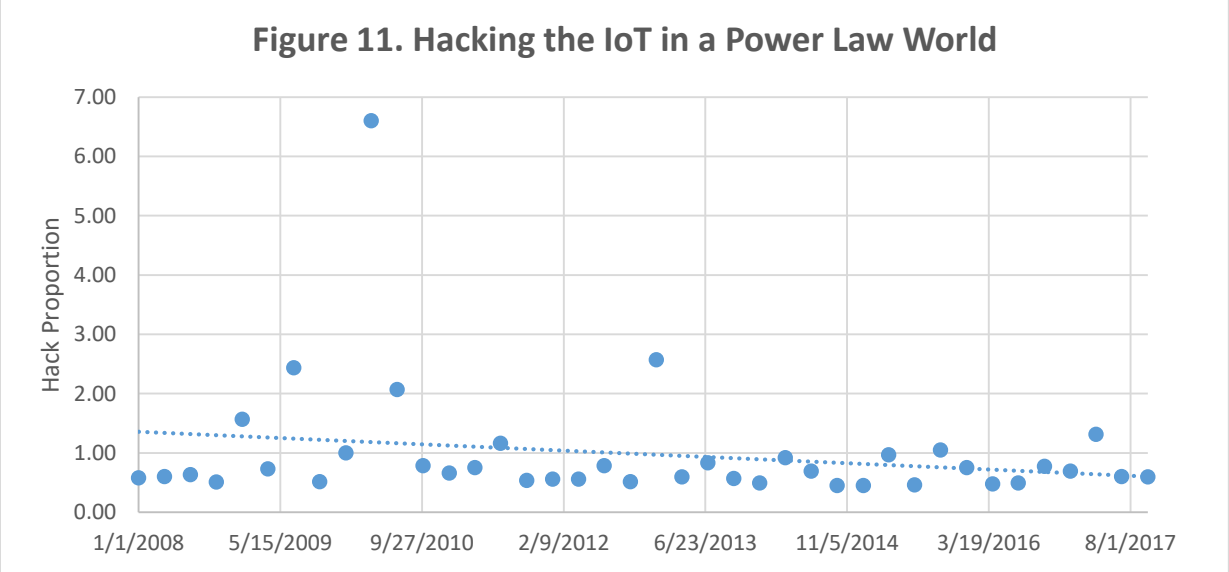

In a one-run simulation, Simpson's Paradox does indeed emerge in a world where cybersecurity is governed by power laws instead of normally distributed curves. Yet, as was the case with the graphical data in the Gaussian world presented above, a single iteration of simulated data reveals very little about what might be going on in more general terms. Indeed, because means are so prone to fluctuations in power law distributions, the graphically presented numbers could be very far off from the more general trend.

To help isolate for whether Simpson's Paradox might emerge over many iterations of the data, I again ran Monte Carlo simulation to capture the effect of 1,000 iterations of the data. Of course, the warning about unstable averages in power law distributions cannot be stressed enough, so even the simulated trend over 1,000 iterations should be considered more suspect than what was found in the Gaussian world.

Table 6 presents the results of the 1,000 simulated iterations. Even over this larger swath of far noisier data, the possibility of Simpson's Paradox emerging to cloud our view of cybersecurity remains a distinct possibility. Over the 27 years of the sample, the proportion of savvy users, naïve users and IoT devices that are hacked all declined, while the aggregate trend significantly increased. Naïve users, for instance, are 14.8 percentage points less likely to be hacked in the 2011 to 2017 period when compared to the 1990 to 1996 period. In contrast, the overall aggregate trend sharply increases over this window of time, rising by some 21.3 percentage points. Simpson's Paradox, the Yule Simpson effect or aggregation bias stands in plain sight as a cloud over our real view of what is happening in cybersecurity.

Table 6. Simpson's Paradox Over 1,000 Iterations of the Data (Power Law World)

	Hack Proportion (1990-1996)	Hack Proportion (1997-2003)	Hack Proportion (2004-2010)	Hack Proportion (2011-2017)	Over Sample Percentage Point Change
Savvy Users	57.8%	53.8%	50.9%	49.0%	-8.8%
Naïve Users	80%	73.3%	69.1%	65.2%	-14.8%
IoT	N/A	N/A	89.8%	84.7%	-5.1%
Aggregate Trend	59.9%	59.3%	67.4%	81.2%	21.3%

Conclusion: The Measurement and Policy Implications of a Simpson's Paradox in Cybersecurity

From many different angles, it looks like the state of cybersecurity is getting a lot worse over time. Trends in hacks, data breaches and observed web-based attacks all point in a worsening direction. But, if the Yule-Simpson effect is at work, these aggregate trends can actually be based upon underlying trends that point in the opposite direction. Sometimes three rights really can make a wrong.

Using a series of Monte Carlo simulation encompassing a total of 2,002 iterations of data, I showed that Simpson's Paradox can easily emerge in cyberspace in both a normally distributed world and a realm where cybersecurity is prone to extremes (power laws). The results of these simulations suggest that users and devices might actually be becoming less prone to being hacked over time, even as the aggregate trends show everything getting much worse. As the simulated data shows us, this vexing mathematical quirk which can so easily confuse policymakers and analysts alike emerges in cyberspace if three plausible conditions are met:

- First, there are definably distinct groups of hackable points online;
- Second, these hackable points have different propensities towards being hacked;
- And, finally, the growth of these groups over time is uneven, with the most susceptible category being the fastest growing segment.

When these three readily satisfiable conditions obtain, the result is a version of Simpson's Paradox, where negative aggregate trends mask positive trends across all subsamples. The plausible operation of the Yule-Simpson effect in cyberspace has significant implications for both the measurement of cybersecurity trends and formulation of public policy.

On the measurement front, the possibility of a Yule-Simpson effect in cybersecurity should caution those involved in the collection and analysis of security incident data away from looking only at overall trends as if they unambiguously represent the actual level of security online. The Internet ecosystem is nothing if not dynamic and its ever changing (and ever expanding) base means that plotting over time trends can be inherently fraught with statistical perils. As demonstrated here through the simulated data, the proliferation of relatively insecure IoT devices will make things seem far worse online than they were before. But that does not necessarily mean that users (and the designers of IoT devices, for that matter) are not getting better at preventing unwanted breaches. To accommodate for the easy occurrence of Simpson's Paradox, analysts and data scientists need to collect data that is stratified along relevant lines.

Luckily, the internal mathematical logic of Simpson's Paradox provides some clues as to where we should look. Generic categories do not matter. In order for Simpson's Paradox to emerge, the groups themselves need to be correlated with both the outcome (some form of security incidents) and the independent variable of interest (growth of the network over time, for example). Recognizing this simple requirement helps to immediately narrow the field of potential categories. Despite the 'collect all' mantra of the Big Data age, collecting information still comes with a cost and using theory as a guide can help to reduce the burden.

On the policy front, the possibility of a Yule-Simpson effect in cyberspace poses challenges for those who advocate for extensive policy reform in the face of worsening aggregate trends. A lot hinges here upon the plausibility of the three conditions highlighted above. If the conditions are deemed to be implausible, then the demonstration of the Yule-Simpson effect in cyberspace with simulated data is tenuous. However, if the idea that there are users or end points online that can be categorized, have different hack propensities and grow in roughly inverse relation to their level of security, then Simpson's Paradox is likely, as shown here.

This puts policymakers in a tough spot. Certainly, there are genuine security problems online, ranging from data breaches, to ransomware, hacked cars and even compromised wearable technologies.[39] Additionally, the media has sensationalized data breaches, hacks and IT security to a significant degree and people routinely express concern that they or their national institutions will be the target of a breach or malicious online activity.[40] This combination of real problems and media-hyped sensationalism makes it difficult for policymakers to sit on their hands. Many might want to act and to act decisively and disjunctively, changing everything possible just to be seen to be keeping people and devices safe online.

The trouble is, as was the case with massive educational reform during the 20th century in America, the best policies might be incremental tweaks, modifications and reforms rather than massive overhauls because things might actually be getting better over time. A part of the challenge for policymakers is that we do not really know. Policymakers, broadly defined, need

better evidence, which loops back and links the problems of policymaking right back to the problem of measurement and data collection. Researchers and policymakers, in other words, need to work hand in hand.

References

[1] National Commission on Excellence in Education, "A Nation at Risk: The Imperative for Educational Reform," 1983. Accessed at: https://www2.ed.gov/pubs/NatAtRisk/index.html ; C. C. Carson, R. M. Huelskamp and T. D. Woodall, "Perspectives on Education in America: An Annotated Briefing, April 1992," *Journal of Educational Research*, Vol. 86, no. 5, 1993, 259-265, 267-291, 293-297, 299-307, 309-310.

[2] National Commission on Excellence in Education, "A Nation at Risk: The Imperative for Educational Reform," 1983. Accessed at: https://www2.ed.gov/pubs/NatAtRisk/index.html ; C. C. Carson, R. M. Huelskamp and T. D. Woodall, "Perspectives on Education in America: An Annotated Briefing, April 1992," *Journal of Educational Research*, Vol. 86, no. 5, 1993, 268.

[3] For a somewhat contrary view, see Eric Jardine, "Global Cyberspace is Safer than you Think: Real Trends in Cybercrime," *Global Commission on Internet Governance Paper Series*, no. 16, 2015, 1-22. Accessed at: https://www.cigionline.org/sites/default/files/no16_web_1.pdf ; Eric Jardine, "Garbage In, Garbage Out: Assessing the Effectiveness of Remedial Cybersecurity Policy," *IEEE: Security & Privacy* (under review).

[4] Daniel Roberts, "Tom Ridge: Cyber attacks are now worse than physical attacks." *Yahoo! Finance, 2016*. Accessed online at: http://finance.yahoo.com/news/tom-ridge-cybersecurity-attacks-are-now-worse-than-physical-attacks-170426390.html?soc_src=social-sh&soc_trk=tw

[5] E.H. Simpson, "The Interpretation of Interaction in Contingency Tables," *Journal of The Royal Statistical Society*, Series B, Vol. 13, no. 2, 238-241, 1951.

[6] Nassim Nicholas Taleb, 2010, *The Black Swan: The Impact of the Highly Improbable* (New York: Random House Trade Paperbacks).

[7] Privacy Rights Clearinghouse only counts breaches involving financial information and SSN in their count. These counts include all breached records, which can include things like email addresses and other bits of private, yet not financial, information.

[8] For an alternative interpretation of the data using Bayesian models, see Benjamin Edwards, Steven Hofmeyr and Stephanie Forrest, "Hype and Heavy Tails: A Closer Look at Data Breaches," *Journal of Cybersecurity* Vol. 2, no.1 (2016), 3-14.

[9] Kaspersky Labs, "Kaspersky Security Bulletin 2008." Accessed at: https://securelist.com/analysis/kaspersky-security-bulletin/36241/kaspersky-security-bulletin-statistics-2008/ Kaspersky Labs, "Kaspersky Security Bulletin 2009." Accessed at: https://securelist.com/analysis/kaspersky-security-bulletin/36284/kaspersky-security-bulletin-2009-statistics-2009/ Kaspersky Labs, "Kaspersky Security Bulletin 2010." Accessed at: https://securelist.com/analysis/kaspersky-security-bulletin/36345/kaspersky-security-bulletin-2010-statistics-2010/ Kaspersky Labs, "Kaspersky Security Bulletin 2011." Accessed at: https://securelist.com/analysis/kaspersky-security-bulletin/36344/kaspersky-security-bulletin-statistics-2011/ Kaspersky Labs, "Kaspersky Security Bulletin 2012." Accessed at: https://securelist.com/analysis/kaspersky-security-bulletin/36703/kaspersky-security-bulletin-2012-the-overall-statistics-for-2012/ Kaspersky Labs, "Kaspersky Security Bulletin 2013." Accessed at: http://media.kaspersky.com/pdf/ksb_2013_en.pdf Kaspersky Labs, "Kaspersky Security Bulletin 2014." Accessed at: https://securelist.com/files/2014/12/Kaspersky-Security-Bulletin-2014-EN.pdf Kaspersky Labs, "Kaspersky Security Bulletin 2015." Accessed at: https://securelist.com/analysis/kaspersky-security-bulletin/73038/kaspersky-security-bulletin-2015-overall-statistics-for-2015/ Kaspersky Labs, "Kaspersky Security Bulletin 2016." Accessed at: https://securelist.com/analysis/kaspersky-security-bulletin/72771/kaspersky-security-bulletin-2016-predictions/

[10] For a consolidated overview, see Statista, "Global number of web attacks blocked per day from 2012 to 2015 (in millions)." Accessed at: https://www.statista.com/statistics/494961/web-attacks-blocked-per-day-worldwide/ . See, for example, Symantec, "2013 Internet Security Threat Report." Accessed at: http://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v18_2012_21291018.en-us.pdf Symantec, "2014 Internet Security Threat Report." Accessed at: http://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v19_21291018.en-us.pdf Symantec, "2015 Internet Security Threat Report." Accessed at: https://www.symantec.com/content/en/us/enterprise/other_resources/21347933_GA_RPT-internet-security-threat-report-volume-20-2015.pdf Symantec, "2016 Internet Security Threat Report." Accessed at: https://www.symantec.com/security-center/threat-report

[11] Dinei Florencio and Cormac Herley, "Sex, Lies and Cybercrime Surveys," 2013, "Sex, Lies and Cyber-Crime Surveys," Workshop on Economics of Information Security and Privacy, 1-11.
[12] Nir Kshetri, 2006, "The Simple Economics of Cybercrimes," IEEE: Security & Privacy Vol. 4, no. 1, 33-39. Accessed at: http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=1588823
[13] Matt Matthias Brecht and Thoas Nowey, 2013, "A Closer Look at Information Security Costs," in Rainer Bohme ed. The Economics of Information Security and Privacy, 3-24.
[14] Ross Anderson, Chris Barton, Rainer Bohme, Richard Clayton, Miachel J.G. Van Eeten, Michael Levi, Tyler Moore and Stefan Savage, 2013, "Measuring the Cost of Cybercrime," in Rainer Bohme ed. The Economics of Information Security and Privacy, 265-300.
[15] Ross Anderson, Rainer Bohme, Richard Clayton and Tyler Moore, 2008, "Security Economics and European Policy," Workshop on Economics of Information Security and Privacy. Accessed at: http://www.econinfosec.org/archive/weis2008/papers/MooreSecurity.pdf
[16] Eric Jardine, "Mind the Denominator: Towards a More Effective Measurement System for Cybersecurity," *Working Paper*.
[17] Eric Jardine, "Global Cyberspace is Safer than you Think: Real Trends in Cybercrime," *Global Commission on Internet Governance Paper Series*, no. 16, 2015, 1-22. Accessed at: https://www.cigionline.org/sites/default/files/no16_web_1.pdf
[18] Index of Cybersecurity. Accessed at: http://cybersecurityindex.org/
[19] A.P. David, cited in, Judea Pearl, *Causality: Models, Reasoning, and Inference* (Cambridge: Cambridge University Press, 2009), 78.
[20] P. J. Bickel, E. A. Hammel, J. W. O'Connel, "Sex Bias in Graduate Admissions: Data from Berkeley," *Science, Vol. 187, 1975*. Accessed at: https://en.wikipedia.org/wiki/Simpson's_paradox#cite_note-freedman-12 ; Steve Berman, Leandro DalleMule, Michael Greene and John Lucker, "Simpson's Paradox: A Cautionary Tale in Advanced Analytics," *The Statistics Dictionary*, 2012. Accessed at: https://www.statslife.org.uk/the-statistics-dictionary/2012-simpson-s-paradox-a-cautionary-tale-in-advanced-analytics
[21] C.R. Charig, D.R. Webb, S.R. Payne and J.E.A. Wickham, "Comparison of Treatment of Renal Calculi by Open Surgery, Percutaneous Nephrolithotomy, and Extracorporeal Shockware Lithotripsy," *British Medical Journal*, vol. 292, 1986. Accessed at: https://www.ncbi.nlm.nih.gov/pubmed/3083922
[22] Steven A. Julious and Mark A. Mullee, "Confounding and Simpson's Paradox," *British Medical Journal*, Vol. 309, 1994. Accessed at: http://www.bmj.com/content/309/6967/1480
[23] Jeff Gill, *Essential Mathematics for Political and Social Research*. Cambridge: Cambridge University Press, 2008, 316.
[24] See, for example, the trends in the Index of Cybersecurity. Accessed at: Accessed at: http://cybersecurityindex.org/
[25] Whitney Phillips, *This is Why We Can't Have Nice Things: Mapping the Relationship Between Online Trolling and Mainstream Culture*. Cambridge: MIT Press, 2016.
[26] Gabriela Coleman, *Hacker, Hoaxer, Whistleblower, Spy: The Many Faces of Anonymous*. New York: Verso, 2015.
[27] Brian Krebs, *Spam Nation: The Inside Story of Organized Cybercrime—From Global Epidemic To Your Front Door*. Naperville: Sourcebooks, Inc., 2014.
[28] There is a large literature on the positive effects of knowledge, perceptions of self-efficacy and incentives for proper end user behavior. See, Eric Jardine, "Context Matters: End User Costs When Things Go Wrong and Information Security in the Workplace," *Working Paper*; Chan, Mark. Irene Woon and Atreyi Kankanhalli. 2005. "Perceptions of Information Security in the Workplace: Linking Information Security Climate to Compliant Behavior." International Journal of Information Security and Privacy, vol. 1, no.3, 18-41; Herath, Tejaswini and H.R. Rao. 2009. "Encouraging information Security Behaviors in Organizations: Role of Penalties, Pressures and Perceived Effectiveness." Decision Support Systems, Vol. 47, 154-165; Ifinedo, Princely. 2014. "Information Systems Security Policy Compliance: An Empirical Study of the Effects of Socialisation, Influence, and Cognition." Information & Management, 51, 69-79; Sommestad, Teodor, Jonas Hallberg, Kristoffer Lundholm and Johan Bengtsson. 2014. "Variables Influencing Information Security Policy Compliance." Information Management & Computer Security, vol. 22, no. 1, 42-75.
[29] Robyn Collinge, 2017, "Your Users Might Not be as Tech-Savvy as You Think." *USABilla Blog.* Accessed at: http://blog.usabilla.com/users-might-not-tech-savvy-think/

[30] Mario Ballano Barcena and Candid Wueest, "Insecurity in the Internet of Things," *Symantec*, 2015. Accessed at: https://www.symantec.com/content/dam/symantec/docs/white-papers/insecurity-in-the-internet-of-things-en.pdf

[31] HP Study Reveals 70 Percent of Internet of Things Devices Vulnerable to Attack. Accessed at: http://www8.hp.com/us/en/hp-news/press-release.html?id=1744676#.WSMGKmjysuU

[32] Level 3 Threat Research Labs, "How The Grinch Stole IT," 2016. Accessed at: http://blog.level3.com/security/grinch-stole-iot/

[33] Data for population and Internet penetration was taken from the World Bank and data on the number of IoT devices was taken from Statista Research. See, World Bank Indicators, http://data.worldbank.org/indicator and Statista Research, "Internet of Things (IoT): number of connected devices worldwide from 2012 to 2020 (in billions)." Accessed at: https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/

[34] The linear interpolation equation is first value+(final value-first value)/number of intervening periods.

[35] Eric Jardine, 2017. "Something is Rotten in the State of Denmark: Why the Internet's Advertising Business Model is Broken and What can be Done About It." *First Monday*, vol. 22, no. 7. Accessed at: http://dx.doi.org/10.5210/fm.v22i7.7087; Lada A. Adamic and Bernardo A. Huberman. 2000. "Power-Law Distribution of the World Wide Web." Science, Vol. 287, no. 5461. 2115a; Michalis Faloutsos, Petros Faloutsos and Christos Faloutsos. 1999. "On Power-Law Relationships of the Internet Topology." Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols For Computer Communication. Accessed at: http://www.cs.cmu.edu/~christos/PUBLICATIONS/sigcomm99.pdf ; Johnson, Steve L., Samer Faraj and Srinivas Kudaravalli. 2014. "Emergence of Power Laws in Online Communities: The Role of Social Mechanisms and Preferential Attachment." MIS Quarterly, Vol. 38, no. 3. 795-808.

[36] Dinei Florencio and Cormac Herley, "Sex, Lies and Cybercrime Surveys," 2013, "Sex, Lies and Cyber-Crime Surveys," Workshop on Economics of Information Security and Privacy, 1-11.

[37] Max Boisot and Bill McKelvey. 2007. "Extreme Events, Power Laws, and Adaptation: Towards an Econophysics of Organization." Academy of Management Proceedings; Newman, M.E.J. 2005. "Power Laws, Pareto Distributions and Zipf's Law." Contemporary Physics, Vol. 46, 323-351.

[38] This is basically Taleb's point in The Black Swan. See, Nassim Nicholas Taleb, 2010, *The Black Swan: The Impact of the Highly Improbable* (New York: Random House Trade Paperbacks).

[39] Marc Goodman, *Future Crime: Everything is Connected, Everyone is Vulnerable, and What We Can Do About it*. Toronto: DoubleDay Canada.

[40] Fen Osler Hampson and Eric Jardine, *Look Who's Watching: Surveillance, Treachery and Trust Online*. Waterloo: Centre for International Governance Press, 2016, 156.