

Empirical Analysis of Losses from Business-Email Compromise

Geoffrey Simpson
Tandy School of Computer Science
The University of Tulsa
Tulsa, Oklahoma 74104
Email: geoffrey@utulsa.edu

Tyler Moore
Tandy School of Computer Science
The University of Tulsa
Tulsa, Oklahoma 74104
Email: tyler-moore@utulsa.edu

Abstract—We examine approximately nine months of data on losses from business email compromise (BEC) reported to the FBI’s Internet Crime Complaint Center in 2017. We describe the empirically observed loss distribution. We study differences in the amounts attempted stolen when the attacks were successful or not. We show that money stolen and transmitted internationally is less likely to be recovered. We also find, somewhat surprisingly, that illicit transfers to in-state banks are also more likely to succeed. Finally, we study state-level differences among BEC target selection and asset recovery.

I. INTRODUCTION AND BACKGROUND

Business email compromise (BEC) is a specialized form of cybercrime in which attackers target employees with authority to transfer company funds and dupe them into sending money. BEC fraud has been increasing in popularity, in large part due to the large sums that can be stolen from each victim compared to consumer-targeted phishing. Mimecast reported a growth rate of 80% quarter over quarter [1]. The United States Federal Bureau of Investigation (FBI) has been involved in investigating business email compromise fraud since at least 2013 [2] and has been collecting data along the way. In 2018, the FBI reported that \$1.3 billion was successfully stolen in BEC scams [3]. And this is despite a recent global sweep in which the FBI arrested 74 people around the world perpetrating BEC fraud. In 2019, the FBI reported that BEC losses increased to \$1.7 billion [4].

Despite these alarming figures, much remains unknown, such as about the distribution of losses, whether certain companies (based on geography or business type) are targeted more often, and whether there is any relationship between the amount attempted stolen and criminal success. In this paper, we set to answer these and other questions, using data on all BEC incidents reported to the FBI’s Internet Crime Complaint Center (IC3) during nine months in 2017.

Our work is motivated by the economics of information security literature [5]. This community has shed much light on how attackers operate [6], [7], [8], [9], [10] and the effectiveness of defenses [11], [12], [13], [14].

One lingering challenge, however, has been to obtain accurate empirical measurements of losses associated with various cybercriminal activities. Estimates for aggregate losses have sometimes been tabulated (see, e.g., [15], [16]). While helpful,

in order to be able to manage cybersecurity risks, it is also necessary to determine the *distribution* of loss amounts. Distributions of cybercrime losses matter for two principal reasons. First, it is widely known that the distributions of losses can be skewed, rendering mean and median values unreliable for capturing the magnitude of risk present [17]. Second, models of cybersecurity investment, including the canonical Gordon-Loeb model [18], require a loss distribution function to operate. However, due to the difficulty of obtaining empirically-validated loss distributions, these models typically assume loss distributions selected for mathematical convenience and without regard to how much they correspond to actual losses.

In theory, insurers are in a good position to estimate loss distributions using claims data from cyber insurance products. In practice, however, actuarial data is proprietary and fiercely guarded, so there has been no public reporting of empirically-derived cyber loss distributions by insurers. Furthermore, claims data would only cover the losses paid out by insurers. This leaves out the tails of the distribution: incidents whose cost falls below the deductible or exceeds coverage limits. Given that cyber loss distributions are likely heavy-tailed, these are critical omissions. One recent approach has been to estimate loss distributions from publicly reported prices filed with state Insurance commissioners [19]. While promising, it does not eliminate the need for empirically-derived loss distributions.

Consequently, the primary contribution of this paper is to provide an empirical loss distribution for BEC losses reported to the FBI IC3. It is hoped that this can be used by researchers to validate theoretical models and by practitioners to manage BEC risks in their own organizations. Section III examines both realized and prevented losses, drawing insights about how these differ and examining the relationship between the amount attempted stolen and attack success. Section IV studies how (and whether) BEC frauds vary geographically. Finally, we conclude in Section V.

II. DATASET DESCRIPTION

The FBI IC3 team shared in anonymized form all BEC-related reports received between January 1, 2017 and September 27, 2017. The dataset indicates the time when the incident




Fig. 1. Histogram of successful BEC thefts, grouped by amount stolen. The left graph shows the percentage of total incidents, while the right graph shows how much money is stolen in total.

was reported. It includes the amount attempted stolen, as well as the amount lost and not recovered. It specifies the city, state, and country of the victim and recipient (typically a bank) when available.

We started with 12,651 distinct incidents. We excluded 16 reports where the funds actually lost exceeded the amount attempted stolen. Additionally, we excluded 4,274 reports in which both the amount attempted and successfully stolen were missing or zero.

7,564 reports included information on both the victim and recipient banks. 361 incidents include only victim information, while 436 only include recipient information.

Several limitations of the dataset should be acknowledged. Most notably, the incidents reported to the IC3 are self-reported by victims, which means that there could be selection bias. Larger, more sophisticated companies may be more likely to be aware of IC3 and the need to report. Additionally, organizations based in geographic locations where the FBI has a greater presence may be more likely to report. Finally, some victims may elect to not report the losses for fear of the information causing reputational damage. This effect should be limited, given that company information is not publicly disclosed.

In terms of geographical coverage, we expect that the dataset has much better coverage for BEC fraud targeting US firms than others. The data does include reports on victims internationally, so we only expect the coverage to be comprehensive for US victims and recipient banks. 96% of victims and 84% of identified recipients are based in the US. Additionally, the geographic information (city, state, country) has complete coverage for victims, but the country is missing 24% of the time for recipients. International cities are rarely reported, hence we only consider the country in our analysis.

The data covers nearly nine months of incidents during 2017. We could not influence the time period or duration for which the data was shared. As a result, we cannot draw any conclusions about how the BEC threat has evolved over time.

Despite these limitations, this dataset remains an excellent source due to its vast size and the active, multi-year involvement of the FBI in countering BEC threats. It stands to reason that the IC3 data on BEC fraud is more comprehensive than any other potential source by far.

III. EMPIRICAL ANALYSIS OF LOSSES

Across the 7,925 reported BEC thefts in our dataset, more than \$1.361 billion was attempted stolen. 4,166 (52.5%) of these incidents succeeded in stealing funds, totaling \$480 million. Nonetheless, even more money was put at risk, with approximately \$881 million of funds blocked from being turned over to cybercriminals.

A. Loss distribution for successful BEC attacks

Figure 1 examines the distribution of funds lost to successful BEC frauds. Because the data is so highly skewed, we have grouped the losses into logarithmically-sized buckets. Figure 1 (left) reports the percentage of all incidents with loss amounts in each category. 60% of thefts involve amounts between \$10K–\$100K, with another 20% falling between \$100K–\$1M. Figure 1 (right) shows the fraction of total losses for each category. Despite accounting for most of the incidents, the total amount lost to thefts between \$10K–\$100K is \$97 million, approximately 20% lost overall. Theft between \$100K and \$1 million added up to \$227 million, 47% of the total. Finally, despite only accounting for 1.4% of the incidents, thefts over \$1 million collectively netted \$150 million, 31% of the total haul.

We tested the data to see if it matched any well-known statistical distributions. We computed best fits for Gamma, Weibull, Pareto and Lognormal distributions using maximum likelihood estimation. Unfortunately, these fits were rejected based on the Kolmogorov-Smirnov and Cramer-von Mises test statistics.

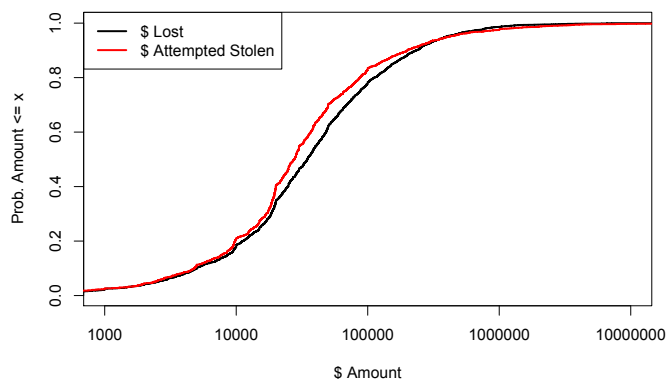


Fig. 2. Cumulative distribution functions (CDFs) of BEC losses, attempted and successful.

B. Recovery of Funds

We now look more closely at this breakdown between funds that are stolen and those that can be protected.

Figure 2 plots the cumulative distribution functions for stolen funds on a logarithmic scale. The red line indicates the amounts attempted to be stolen, while the black line is for the money lost and not recovered. CDFs show the proportion of frauds with losses at or below a specified amount. For example, 78% of the BEC incidents where money was irretrievably lost, the amount stolen was less than \$100,000. By contrast, 83% of the blocked incidents involved amounts less than \$100,000.

Note that for the all but the largest thefts, the CDF for realized losses is to the right of the recovered losses. This means that successful BEC attacks tend to get away with *more* money than when the fraud is blocked. This is reflected in the median amounts lost: \$35K for successful attacks and \$27.5K when blocked. However, there is an important caveat: beginning at 94%, the blocked BEC amounts exceed those of the successful frauds. This means that the 6% biggest blocked BEC frauds are substantially bigger than the 6% biggest successful frauds. This adds up, since the total amount blocked is much greater than the total amount lost (\$881 million vs. \$479 million). The upshot is that blocking these most egregious thefts adds up to the biggest savings overall. (Of course, the large magnitude of attempted theft could very well explain why they are found out and stopped.)

Looking back at Figure 1 (right), we can see that around half of the total lost funds are tied to mid-sized losses between \$100K–\$1M. Figure 3 plots the same histograms for blocked transfers. In comparison, just 20% of total blocked funds transfers fall in the \$100K–\$1M range, with over two thirds of the total blocked accounted for by attempted thefts of over \$1 million. This is particularly striking, given that they comprise just 2.2% of blocked transfers.

It is significant that the distribution of successful losses is so different than the distribution of blocked losses. It demonstrates that one cannot rely exclusively on a loss distribution for realized losses when computing a cost-benefit analysis of investing in security controls. When defensive

\$ Attempted	% Incidents Successful	Statistically Significant?	Standardized Residuals
Overall	52.9%		
≤ \$10K	48.8%	(–)	–2.22
\$10K < \$100K	51.9%		–1.15
\$100K < \$1M	62.3%	(+)	4.65
> \$1M	45.4%		0.160

TABLE I

PROPORTION OF ATTEMPTED ATTACKS THAT SUCCEEDED IN STEALING MONEY, GROUPED BY AMOUNT ATTEMPTED STOLEN.

countermeasures successfully thwart attacks, the benefits can actually exceed the realized losses that others experience.

We now dig a bit deeper into the question of whether the amount of money that is attempted stolen could affect fraud success. Table I compares the proportion of BEC frauds succeed in stealing funds grouped by size. Overall, 53% of reported BEC frauds get away with the money. However, the success rate is actually lower for small frauds under \$10K. Just 48% of these smaller frauds succeed; this difference is statistically significant according to a χ^2 test¹. By contrast, larger frauds between \$100K and \$1M succeed most often, registering at 62% of the time. The very largest frauds succeed less often, but the difference is not statistically significant.

Finally, we now consider whether the location of the bank receiving the fraudulent payment affects success. In 18% of reported cases, the victim is located in a different country than the bank receiving the fraudulent transfer. When this happens, funds are lost 77% of the time. When the recipient bank is located in the same country as the victim, the criminals are less likely to successfully abscond with the funds (doing so only 56% of the time). These differences in percentages are statistically significant according to a χ^2 test.

What about when the victim and recipient bank are located in the same state? This happens in around 9.8% of domestic cases. Surprisingly, these transfers are much less likely to be stopped. 62.2% of intrastate BEC frauds succeed, compared to 55% of interstate transfers. Again, this difference is statistically significant.

IV. DOES BEC VARY BY GEOGRAPHY?

Cybercrime reports regularly break down attacks or losses by geographical boundaries such as countries or states [4]. It is a natural and straightforward exercise to do so, but it is not obvious why we would even expect there to be significant differences once the data has been normalized. Large-scale cybercrimes like BEC cast a wide net and do not typically tailor their methods to particular regional characteristics.

But how should that normalization take place? It turns out that for BEC, like many cybercrimes, there are several ways data could be normalized. We compare the possibilities in Section IV-A. Next, we study normalized BEC victim and recipient rates across states in Section IV-B, followed by success rates in different states in Section IV-C.

¹Throughout the paper, any time we report a test being statistically significant, it is significant with at least 95% confidence.




Fig. 3. Histogram of blocked BEC thefts, grouped by amount attempted stolen. The left graph shows the percentage of total incidents, while the right graph shows how much money would have been stolen in total had the money not been recovered.

A. What is the best way to normalize BEC rates?

It is interesting to compare BEC fraud rates by state. However, it is clear that direct comparison between states of fraud activity is not particularly meaningful, since there can be vast differences in population and economic activity in those states. The most straightforward approach is to normalize by population. But is that the best approach for BEC fraud rates? After all, these frauds target firms, not consumers. Moreover, financial institutions receive, process and (hopefully) block fraudulent payments. Thus we now investigate more closely how to appropriately normalize BEC frauds across geographic jurisdictions.

We consider four candidate normalization variables:

- 1) Population
- 2) # Public Companies Headquartered in State
- 3) # Banks Headquartered in State
- 4) Total Bank Assets

We do not know the types of companies in the dataset, and they are probably a mix of private and public companies, ranging in size. We collect information on the location of public company headquarters in the United States. We compiled the data manually by visiting the Nasdaq website and searching for public companies by state. It is hoped that this figure represents a good proxy of overall business activity in states, and therefore, a reasonable measure of prospective BEC targets.

In terms of BEC fraud recipients, we expect that to depend on the banking activity present in each state. We therefore gather reports of the number of banking institutions located in each state, as well as the total financial assets of these banks, from the FDIC [20]. We use data reported for 2017 in order to be consistent with the timing of the BEC data. Because asset size is so skewed, we perform a log transformation.

Table II reports the correlation matrix for these candidate normalization variables. We see that, unsurprisingly, they are highly correlated. Population and public companies are most

	Population	# Public Cos.	# Banks	log(Bank Assets)
Population	1.00	0.85	0.46	0.48
# Public Cos.	0.85	1.00	0.42	0.49
# Banks	0.46	0.42	1.00	0.36
log(Bank Assets)	0.48	0.49	0.36	1.00

TABLE II
CORRELATION MATRIX FOR CANDIDATE NORMALIZING VARIABLES.

Indep. Variable	Dependent Variable	
	# Victims	# Recipients
Population	0.92	0.84
# Public Cos.	0.73	0.55
# Banks	0.11	0.20
log(Bank Assets)	0.16	0.20


TABLE III
 R^2 FOR VARIOUS LINEAR REGRESSIONS USING DIFFERENT INDEPENDENT AND DEPENDENT VARIABLES.

closely correlated (at $r = 0.85$). The correlation between the banking variables and the others are lower, but still rather high.

As a result, we cannot reliably include more than one normalization variable at a time in any linear regressions. Instead, we compute regressions using each of the four candidate variables as the sole independent variable, with the number of victims and recipients as dependent variables. We exclude the full tables with coefficients for brevity, noting that in all regressions the coefficients for the independent variables were positive and significant. Instead, Table III reports the R^2 for each of the regressions. This indicates the amount of variance between the number of BEC victims or recipients per state that can be explained by the independent variable.

Population explains the most variance by far: 92% of the variance of victims per state and 84% of the variance of recipients. The number of public companies headquartered in a state explains 73% of the variance of the BEC victims in that state, suggesting that it is a good proxy indicator of how many

BEC Victims per 1M Residents



BEC Recipients per 1M Residents




Fig. 4. Heatmaps of normalized BEC fraud rates for victims (left) and recipients (right), normalized by state population.

targets there are (though not as good as population). The link to recipients is weaker, however. As expected, the number of banks and size of assets is not nearly as highly correlated with the number of bank victims. Surprisingly, though, both bank figures do not do much better when explaining the variance in the number of BEC recipients by state.

Our conclusion from this analysis is that, of the indicators presented, normalizing by the most straightforward measure, population, is best.

B. Victims and recipients by state

Table IV reports the number of victims and recipients broken down by state. It normalizes all figures by the state's population, reporting the number and amount lost per million residents. For example, Washington DC has the most number of victim per capita, at 62 victims per million population. DC is an outlier for many reasons, but the next most common are Colorado, Nevada, Guam, California, Massachusetts and Florida, each registering between 30–40 victims per million residents. In terms of financial losses, Arizona tops the table at \$5.7 million lost per million residents, followed by Colorado at \$3.3 million.

But what about those receiving BEC funds? Organizations in Delaware, DC, Maryland and Georgia receive the most stolen funds, from 33–39 events per million residents. In terms of the financial amount received, South Dakota is far and away the biggest, with \$32 million received per million residents. This reflects the outsize number of large financial institutions present despite its low population. Notably, while there is some relationship between the number of BEC victims and recipients, the correlation is relatively weak at $r = 0.31$.

Figure 4 plots the number of BEC victims and recipients for each state, normalized by population. It is clear from the graphs that the rate of BEC victimization in a state has little bearing on where the recipients of such fraud are located.

Figure 5 provides a series of scatter plots comparing the number of victims and recipients in each state to its population. Also included is the best-fit line obtained from running a robust linear regression using a bisquare weighting function to account for outliers.

The leftmost graph plots the number of victims in each state against its population. The graph is dominated by the states with very large populations, so the graph next to it zooms in and looks at only those states with populations under 12 million. The significance of the best-fit line is that states above the line have more BEC victims than would be expected for their populations, while those below the line have fewer BEC victims per state than expected. So while California has the most people, it has even more BEC victims than the model would predict. Other overrepresented states include Florida, New York, Colorado and Massachusetts. By contrast, Puerto Rico, Oklahoma, Tennessee and Ohio have fewer BEC victims than their populations would suggest.

The right graphs in blue in Figure 5 compare the number of BEC recipients to state populations. Visually, we can see greater dispersion around the best-fit line, which is reflected in the lower R^2 value reported in Table III. Here, California is much closer to the expected number of recipients than it was for victims. Florida remains overrepresented, while Texas is significantly overrepresented as recipients of BEC frauds but not for victims. Puerto Rico is even more underrepresented as a recipient than for victims. And Washington is overrepresented among victims but underrepresented among recipients.

Nonetheless, these differences across jurisdictions are relatively minor, and this is not particularly surprising given that BEC targets companies in all US jurisdictions with similar tactics.




Fig. 5. Scatter plot of BEC victims per state (left, in black) and recipients per state (right, in blue).

C. Does fraud success vary by state?

From the analysis in Section III-B we know that slightly more than half of attempted BEC frauds reported to the FBI succeed in stealing funds. We now consider whether that success rate could vary by state.

We perform two tests. First, we compare the proportion of BEC frauds that succeed across all states by inspecting the victim reports. While there is considerable variation, none of these differences are statistically significant according to a χ^2 test. Hence, we cannot conclude that criminals are more successful stealing from victims in any state.

For the second test, we check the proportion of BEC recipients that fail to block stolen funds in each state. Here, a few states were statistically significantly different than the rest. Banks in New York that received incoming fraudulent payments let the money through 67% of the time. In Kentucky, by contrast, recipients blocked the payments 72% of the time, and Michigan recipients blocked them 58% of the time. All other states were statistically indistinguishable from the average blocking rate of 42%.

V. CONCLUDING REMARKS

We have set out to explore in greater detail the monetary losses associated with business-email compromise. In terms of cybercriminal revenue, BEC is one of the biggest threats today. Its seriousness is magnified by the fact that individual frauds regularly net hundreds of thousand to millions of dollars. While headline figures have been reported previously, in this paper we examine these losses in much greater detail.

We have reported on what the distribution of losses look like. We have shown that the distribution of blocked frauds differs substantially from that of successful attacks. Most notably, the bulk of money stolen in successful attacks comes from thefts in the hundreds of thousands of dollars, whereas the bulk of blocked funds can be tied to failed attempts to steal more than \$1 million in one go. We have found that, surprisingly, small thefts succeed less often, while thefts in the hundreds of thousands are most likely to succeed. We presented evidence that, as expected, transfers to international banks succeed much more often than domestic thefts. But we also found that among domestic transfers, frauds where the

victim and recipient are in the same state succeed more often than those where the recipient is out of state.

We also investigated state-level differences in fraud, finding that normalizing by population is the most reliable way to explain differences in BEC activity. We also showed that, controlling for population, some states do indeed experience more fraud than others. Nonetheless, our biggest takeaway from the geographic analysis is that state-level differences in BEC fraud rates are not particularly informative and should perhaps not be emphasized in reporting.

To those writing future reports of cybercriminal activities, be it on BEC or other threats, we make a few recommendations informed by our own experience. First, always normalize when comparing population groups. Second, carefully consider the options for normalization. While we evaluated factors such as firm and financial institution activity, we ultimately found population to be the best way to normalize. Third, consider whether it even makes sense to break down cybercriminal activity by geography, when so much cybercrime is not geographically bounded.

A number of open questions remain. Perhaps the most tantalizing would be to study the effects of the FBI's 2018 establishment of the Recovery Asset Team [4]. It would be interesting to determine whether and by how much the success rates of frauds fell in response. This would of course require access to more recent data than has been analyzed here. Moreover, access to data over a longer period would enable us to answer questions about whether or not the effects observed in this paper in fact change over time.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the FBI IC3 team for sharing the anonymized data, without which this research would not have been possible. The authors are supported by US National Science Foundation Award No.1652610.

REFERENCES

- [1] Mimecast, "Security risk assessment," 2018, <https://www.mimecast.com/globalassets/documents/infographics/esra-infographic---august-2018.pdf>.
- [2] F. B. of Investigation, "Business e-mail compromise: Cyber-enabled financial fraud on the rise globally," 2017, <https://www.fbi.gov/news/stories/business-e-mail-compromise-on-the-rise>.

State	Victims		Recipients	
	# per M	\$ per M	# per M	\$ per M
AK	25.7	\$121021	8.1	\$382995
AL	13.3	\$445165	15.2	\$601629
AR	12.6	\$273478	7.7	\$454319
AZ	25.1	\$5729398	11.0	\$1058020
CA	34.0	\$1907551	15.2	\$2853784
CO	39.2	\$3322638	9.3	\$784916
CT	28.4	\$1515624	9.8	\$794352
DC	62.0	\$1990514	36.0	\$5171199
DE	11.4	\$408731	39.5	\$7561106
FL	30.9	\$2469010	28.0	\$6565239
GA	20.3	\$1142461	33.4	\$5463447
GU	36.2	\$1953534	6.0	\$41275
HI	20.3	\$334000	6.3	\$228619
IA	10.2	\$162473	7.0	\$310477
ID	20.4	\$1448962	5.2	\$163144
IL	22.8	\$884590	13.6	\$2472957
IN	16.0	\$721268	8.5	\$565188
KS	18.9	\$305047	8.9	\$460685
KY	13.9	\$643574	7.4	\$520389
LA	19.4	\$1005447	9.6	\$767007
MA	32.4	\$2601755	9.9	\$3577557
MD	23.8	\$649654	34.5	\$3638900
ME	16.5	\$328746	16.5	\$523009
MI	17.5	\$834866	8.0	\$544765
MN	20.4	\$591936	15.1	\$1577830
MO	21.3	\$738415	9.0	\$505865
MS	9.4	\$544027	12.4	\$249735
MT	18.1	\$119325	6.7	\$140098
NC	17.0	\$634352	15.2	\$1811947
ND	18.5	\$1731311	7.9	\$133524
NE	11.5	\$563605	6.8	\$583761
NH	25.3	\$378939	7.4	\$527069
NJ	24.9	\$2223637	15.9	\$1687994
NM	15.3	\$479622	9.6	\$6097342
NV	36.4	\$2121793	16.7	\$937181
NY	25.2	\$1422301	20.5	\$2715084
OH	16.8	\$834689	11.2	\$1057085
OK	13.0	\$869785	10.7	\$608823
OR	24.1	\$409206	12.8	\$1638955
PA	18.0	\$822101	10.0	\$1079462
PR	4.8	\$84649	0.3	\$22876
RI	25.5	\$818176	20.8	\$4194404
SC	16.9	\$998998	11.9	\$1380742
SD	12.6	\$2052890	18.4	\$31994001
TN	13.8	\$361091	14.3	\$907054
TX	21.1	\$1239396	25.9	\$4877782
UT	25.5	\$1166721	9.7	\$529721
VA	25.9	\$1621335	22.6	\$2335865
VT	25.7	\$260040	6.4	\$165593
WA	27.1	\$2869709	8.4	\$1701612
WI	14.1	\$435248	6.4	\$241877
WV	7.2	\$98791	7.7	\$1043711
WY	20.7	\$2226117	24.2	\$995309

TABLE IV
PER-STATE VICTIM AND RECIPIENT FIGURES, NORMALIZED BY
POPULATION.

- “The role of internet service providers in botnet mitigation: An empirical analysis based on spam data,” OECD Publishing, Tech. Rep., 2010.
- [8] Y. Liu, A. Sarabi, J. Zhang, P. Naghizadeh, M. Karir, M. Bailey, and M. Liu, “Cloudy with a chance of breach: Forecasting cyber security incidents,” in *24th USENIX Security Symposium (USENIX Security 15)*. Washington, D.C.: USENIX Association, Aug. 2015, pp. 1009–1024. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/liu>
- [9] K. Levchenko, A. Pitsillidis, N. Chachra, B. Enright, M. Félegyházi, C. Griery, T. Halvorson, C. Kanich, C. Kreibich, H. Liu, D. McCoy, N. Weaver, V. Paxson, G. M. Voelker, and S. Savage, “Click trajectories: End-to-end analysis of the spam value chain,” in *IEEE Symposium on Security and Privacy*, Oakland, CA, May 2011, pp. 431–446.
- [10] N. Leontiadis, T. Moore, and N. Christin, “A nearly four-year longitudinal study of search-engine poisoning,” in *Proceedings of ACM CCS 2014*, Scottsdale, AZ, Nov. 2014.
- [11] T. Moore and R. Clayton, “Examining the impact of website take-down on phishing,” in *APWG eCrime Researchers Summit*, ser. ACM International Conference Proceeding Series, L. F. Cranor, Ed., vol. 269. ACM, 2007, pp. 1–13. [Online]. Available: <https://tylermoore.utulsa.edu/ecrime07.pdf>
- [12] —, “The consequence of non-cooperation in the fight against phishing,” in *APWG eCrime Researchers Summit*. IEEE, 2008, pp. 1–14. [Online]. Available: <https://tylermoore.utulsa.edu/ecrime08.pdf>
- [13] H. Liu, K. Levchenko, M. Félegyházi, C. Kreibich, G. Maier, G. M. Voelker, and S. Savage, “On the effects of registrar-level intervention,” in *Proceedings of the 4th USENIX Conference on Large-scale Exploits and Emergent Threats*, ser. LEET’11. Berkeley, CA, USA: USENIX Association, 2011, pp. 5–5. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1972441.1972448>
- [14] L. Metcalf and J. M. Spring, “Everything you wanted to know about blacklists but were afraid to ask,” Software Engineering Institute - Carnegie Mellon University, Tech. Rep., 2013.
- [15] R. Anderson, C. Barton, R. Böhme, R. Clayton, M. van Eeten, M. Levi, T. Moore, and S. Savage, “Measuring the cost of cybercrime,” in *11th Workshop on the Economics of Information Security (WEIS)*, 2012. [Online]. Available: <https://tylermoore.utulsa.edu/weis12.pdf>
- [16] R. Anderson, C. Barton, R. Böhme, R. Clayton, C. Ganan, T. Grasso, M. Levi, T. Moore, S. Savage, and M. Vasek, “Measuring the changing cost of cybercrime,” in *18th Workshop on the Economics of Information Security (WEIS)*, 2019. [Online]. Available: <https://tylermoore.utulsa.edu/weis19cost.pdf>
- [17] D. Florencio and C. Herley, “Sex, lies and cyber-crime surveys,” in *Workshop on the Economics of Information Security*, June 2011. [Online]. Available: <https://www.microsoft.com/en-us/research/publication/sex-lies-and-cyber-crime-surveys/>
- [18] L. Gordon and M. Loeb, “The economics of information security investment,” *ACM Transactions on Information and System Security*, vol. 5, no. 4, pp. 438–457, Nov. 2002.
- [19] D. W. Woods, T. Moore, and A. C. Simpson, “The county fair cyber loss distribution: Drawing inferences from insurance prices,” in *18th Workshop on the Economics of Information Security (WEIS)*, 2019. [Online]. Available: <https://tylermoore.utulsa.edu/weis19loss.pdf>
- [20] F. D. I. Corporation, “Fdic state profiles,” 2017, <https://www.fdic.gov/bank/analytical/stateprofile/>.
- [3] —, “2018 Internet Crime Report,” 2019, https://pdf.ic3.gov/2018_IC3Report.pdf.
- [4] —, “2019 Internet Crime Report,” 2020, https://pdf.ic3.gov/2019_IC3Report.pdf.
- [5] R. Anderson and T. Moore, “The economics of information security,” *Science*, vol. 314, no. 5799, pp. 610–613, 2006. [Online]. Available: <https://tylermoore.utulsa.edu/science-econ.pdf>
- [6] T. Moore, R. Clayton, and R. Anderson, “The economics of online crime,” *Journal of Economic Perspectives*, vol. 23, no. 3, pp. 3–20, Summer 2009. [Online]. Available: <https://tylermoore.utulsa.edu/jep09.pdf>
- [7] M. van Eeten, J. M. Bauer, H. Asghari, S. Tabatabaie, and D. Rand,