



Quantifying the Pressure of Legal Risks on Third-party Vulnerability Research

Alexander Gamero-Garrido, Stefan Savage, Kirill Levchenko, and Alex C. Snoeren
University of California, San Diego

ABSTRACT

Product vendors and vulnerability researchers work with the same underlying artifacts, but can be motivated by goals that are distinct and, at times, disjoint. This potential for conflict, coupled with the legal instruments available to product vendors (e.g., EULAs, DMCA, CFAA, etc.) drive a broad concern that there are “chilling effects” that dissuade vulnerability researchers from vigorously evaluating product security. Indeed, there are well-known examples of legal action taken against individual researchers. However, these are inherently anecdotal in nature and skeptics of the chilling-effects hypothesis argue that there is no systematic evidence to justify such concerns. This paper is motivated by precisely this tussle. We present some of the first work to address this issue on a quantitative and empirical footing, illuminating the sentiments of both product vendors and vulnerability researchers. First, we canvas a range of product companies for explicit permission to conduct security assessments and thus characterize the degree to which the broad software vendor community is supportive of vulnerability research activities and how this varies based on the nature of the researcher. Second, we conduct an online sentiment survey of vulnerability researchers to understand the extent to which they have abstract concerns or concrete experience with legal threats and the extent to which this mindset shapes their choices.

KEYWORDS

vulnerability; public policy; copyright

1 INTRODUCTION

Software of any complexity is invariably imperfect, riddled with design flaws or deviations from the designers’ intent that are capable of producing unexpected side-effects. While most of these bugs are benign, a subset is of particular concern because they allow an adversary to violate key security properties that would otherwise be assured. Unfortunately, such security vulnerabilities rarely manifest in the absence of specific adversarial inputs and

thus can be extremely challenging to find. Thus, even while a range of development practices (e.g., Microsoft SDL) are thought to reduce their prevalence, it is widely understood that all software ships with security vulnerabilities present. Indeed, much of modern operational security practice today revolves around managing the problems created when these flaws are identified after a product has been deployed.

How security vulnerabilities are discovered is a key aspect of this situation. While many such vulnerabilities are found by the developer, or groups under contract to them, even cursory analysis of empirical vulnerability data shows that the vast majority of critical vulnerabilities are identified by third-party vulnerability researchers who audit software of their own accord and for a variety of reasons. Indeed, this is not surprising as there are far more such independent vulnerability researchers than any one software developer can possibly employ. However, the role of such researchers occupies a conflicted policy space. On the one hand, it is clear that this set of independent research communities is key to both identifying significant flaws in deployed software and creating an incentive for developers to fix them. At the same time, the costs of such discoveries (particularly when uncoordinated with the software developer), in both labor and brand damage, has the potential to create an adversarial relationship between developers and the research community.

Moreover, under existing U.S. law, software developers have a range of legal theories under which they may challenge third-party security research. These include violations of explicit contractual terms in so-called shrink/click-wrap contracts (e.g., it is increasingly common for such documents to explicitly forbid reverse-engineering), violations of the Computer Fraud and Abuse Act (CFAA) covering unauthorized access to computer systems (of particular concern for products with components hosted on third-party services or products which are leased as a service and not purchased outright), violations of electronic communications privacy laws (e.g., the Wiretap Act, Pen/Trap Statute and the Electronic Communications Privacy Act (ECPA) as they apply to research methods that involve intercepting messages sent from a product), as well as more generic claims of libel, trade secret misappropriation, copyright infringement, etc.

Perhaps the best-known example of these tensions arises due to the Digital Millennium Copyright Act (DMCA) whose anti-circumvention requirements were originally designed to protect media publishers against unauthorized copyright violations, but have been read to encompass a range of software protection mechanisms typically encountered when performing security audits. The

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS '17, October 30–November 3, 2017, Dallas, TX, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-4946-8/17/10...\$15.00

<https://doi.org/10.1145/3133956.3134047>

DMCA provides a private course of action for copyright holders to litigate both injunctive relief and monetary damages against violators.¹ Threats under the DMCA have been documented for several high-profile vulnerability research efforts, most famously in 2000 against a group of Princeton researchers forced to withdraw an accepted paper from publication [31].

During periodic rule-making adjustments mandated by the DMCA statute, a variety of industry groups have resisted efforts to create additional safe harbors for security researchers from this legal recourse [34]. This in turn has led to a claim that the DMCA [13, 17], and the threat of its use, has a “chilling effect” on security research. Advocates of reform argue that the ultimate harm is to consumers who are denied an independent security assessment and the resulting improvements in software security. Opponents of reforming the anti-circumvention clause point out that this claim is only documented via anecdote and not by any systematic assessment.

This narrative is not unique to the DMCA, and variants of the “chilling-effects” hypothesis have been proposed across the range of potential legal risks encountered by security researchers, with opposing voices arguing that there is no compelling data to justify any change in policy. Indeed, to date we are unaware of any grounded attempt to quantify the legal risks faced by security researchers that would place the role played by such forces on an evidence-based footing. Within this milieu, there are two related key questions. First, the extent to which modern product companies reserve or assert their legal rights to limit, control or dissuade vulnerability researchers from independently assessing the security of their products. And second, the extent to which vulnerability researchers factor the possibility of adverse legal action when deciding whether to analyze a particular product.

Our work provides a first step to place these questions on an empirical, quantifiable basis. To elucidate the legal posture of product companies we conducted an empirical study of 75 companies (spanning a range of company sizes and product categories) in which each were contacted by security researchers seeking prior approval for independent security assessments of their products. By varying the nature of the request (e.g., whether or not the DMCA is explicitly mentioned) and pedigree of researcher making the request (e.g., academic vs. independent researchers) we sought to tease apart dependent factors that we hypothesized might impact their responses. To understand the role played by researcher experience and perception, we surveyed over 100 vulnerability researchers and evaluated their predisposition to view legal concerns as a key factor in pursuing a given research target and the extent to which this sentiment was driven purely by abstract concerns or whether they had experienced concrete legal threats in practice.

Two key results stand out from our study:

- While some product manufacturers have embraced the role of third-party vulnerability research and provide explicit or implicit consent to conduct such work (either unrestricted or with time-limited coordinated disclosure policy), most are loathe to surrender legal recourse and either are unwilling

to engage on questions of permission or impose significant restrictions on doing so. Moreover, we find a significant difference in the responsiveness afforded to academic vs. independent security researchers.

- Legal concerns are a significant concern for many vulnerability researchers and almost a quarter report having experienced legal threats or action in the course of their research.

The remainder of the paper presents the methodology, data and analysis for both of our measurement instruments, and discussion regarding potential implications of our findings.

2 STUDY METHODOLOGY

One of the ways researchers can protect themselves from the threat of legal action is to secure the consent of the companies whose products they plan to investigate. Although consent does not eliminate litigation risk, it puts a researcher in a dramatically stronger position should a company take legal action. However, anecdotally few researchers seem to take advantage of this potential safe harbor – perhaps fearing that a negative response to such a request could jeopardize their research. Thus, it is not well-understood if companies are amenable to working with researchers in this way or if they prefer to reserve their legal options (*i.e.*, and thus insist on researchers absorbing all such risk). To explore this question empirically, we worked with four security researchers to request explicit authorization from 75 different companies to conduct security evaluations on their products.² The remainder of this section describes our methodology followed by an analysis of the responses in Section 3.

2.1 Researcher selection

We hypothesized that the reputation and affiliation of the researcher making the request might influence a company’s willingness to grant permission. To test our hypothesis, we approached four security-vulnerability experts: two academics and two independent researchers. This experimental was based on our anecdotal understanding that academic researchers benefit from the imprimatur of their host university, and the associated public “optics” associated with their public mission (in addition to the significant resources of a dedicated university counsel).

By contrast, individual independent security researchers may have little or no institutional support and may be far more fragile to legal threats as a result. Further, different researchers have established reputations for how they manage vulnerability disclosure and this may in turn modulate the apprehension potentially felt by companies whose products are being scrutinized. To this end, our group of academic researchers includes one senior faculty member (with tenure), and one junior (tenure-track) faculty member at a different institution. We also recruited two independent researchers: one based in the United States, and another based in the European Union.

¹The Librarian of Congress, who has statutory authority to grant certain exceptions to the DMCA anti-circumvention clause, has recently granted a limited exemption to the “prohibition against circumvention of technological measures controlling access to copyrighted works” [27] for particular security uses.

²We submitted our protocol to our institutional review board (IRB) in advance of this study and they declared it to not be human subjects research because, among other reasons, it focused on organizational responses and not on individuals. However, we were still careful to minimize the overhead on the organizations being evaluated – limiting interactions to written responses and we have chosen not to name the individual companies to avoid any reputational harm.

Recruiting researchers proved to be challenging. We approached two academics and one independent researcher that were initially willing to engage with this project, but then later decided to withdraw from it. Two of the researchers feared company retribution, given that they personally or their staff interact with product manufacturers frequently, and some of their projects receives private funding from such companies. Further, participation in this study offered limited value to them, while potentially endangering their prospects to do actual vulnerability research on some of these products, should the company reject their request in this study.

Both of the academic researchers who ultimately assisted with our study were well-known security faculty with over 150 published papers and 35,000 citations between them. However, the two varied in seniority; one received their Ph.D. roughly a decade before the other. The senior faculty member was also involved in a pilot round of the study (using the same methodology) thus accounting for a disparity in the number of companies assigned to each. The two participating independent researchers also came with well-established track records. The U.S.-based researcher has been involved in security research “on behalf of Fortune 500 enterprise security teams,” and the E.U. researcher has had their work featured in the *New York Times* and *Ars Technica*. Both have been speakers in major independent security conferences (e.g. BlackHat).

2.2 Company selection

We chose the consumer electronics and software industry as a starting point since companies there are likely to have set procedures for dealing with random contacts about security vulnerabilities. Additionally, consumer products are particularly salient in policy disputes surrounding the DMCA anti-circumvention provision [13]. To that end, we excluded non-profits and companies that sell primarily to enterprise customers. However, there are fewer manufacturers of consumer-facing products sold in the U.S. fitting our criteria than one might surmise and we only were able to identify roughly 120 such companies using the methodology described below.

We used five sources to identify candidate companies and pruned this set to find those manufacturing consumer-focused products sold in the U.S.

Fortune 1000. Our first source, aimed at tallying large companies, was the extended list of one thousand firms compiled by *Fortune* magazine [6], typically referred to as the *Fortune 500*. The authors rank companies by “total revenues for their respective fiscal years. Included in the survey are companies that are incorporated in the U.S. and operate in the U.S. and file financial statements with a government agency.” These firms are by any metric some of the largest corporations in the nation. Consistent with our industry focus, we contacted companies from the following categories, as defined by *Fortune*: Computer Software, Computers, Office Equipment, Computer Peripherals, Electronics, Electrical Equipment, Network and Other Communication Equipment, Semiconductors and Other Electronic Components

Large retailers. We looked for manufacturers featured on the “Electronics” section of two of the largest U.S. online retailers: *Target.com* and *Amazon.com*. There, we gathered the products listed on the first page of each subcategory, and filtered for those meeting

our criteria. Note that we did not use these products directly, but rather compiled a list of companies and then followed our product-selection protocol described in Section 2.3.

Stock indices. Our sample of large consumer-technology companies includes components of two stock-related lists: an index compiled by S&P Dow Jones Indices (U.S. Technology Index [2]) and an exchange-traded fund by BlackRock (iShares U.S. Technology ETF [8]), the world’s largest investment firm [24].

Press lists. Contrary to our intuition, many popular consumer products are not manufactured by very large companies, but rather by mid-sized firms specializing in a technological niche. In order to expand our sample to include them, we looked at popular press publications and columns specializing in consumer technology. To protect the anonymity of the companies studied we do not share the URLs of the lists we used, but instead we aggregate them by publication and product category below:

- *cnet.com*: Automotive GPS devices, computers and accessories, smart home devices, and multimedia devices.
- *Fast Company*: Computers and accessories, smart home devices, and tablets.
- *MIT Technology Review*: Computers and accessories, multimedia devices, smart home devices, and wearable electronic devices.
- *PCMag.com*: Drones, multimedia devices, networking devices, printers, and smart home devices.
- *TopTenReviews.com*: Computers and accessories, consumer software, multimedia devices, networking devices, and smart home devices.
- *The Wall Street Journal*: Multimedia devices, networking devices, and smart home devices.
- All others³: Computers and accessories, multimedia devices, networking devices, and smart home devices.

Y Combinator. In order to find early-stage startups with credible prospects, we consulted the list of companies launched by well-known incubator Y Combinator, whose alumni “companies have a combined valuation of over \$80 billion.” [19]

Among this group, a secondary goal was to diversify the size and age of companies contacted—from very large and established firms to startups—in the hope of obtaining a more accurate representation of vulnerability-related policies for a wider range of products. We hypothesized that very large companies who manage third-party vulnerability disclosure on a regular basis may tend to understand the complexities of the problem well and have clean processes for handling such questions. Indeed, many such companies employ third-party organizations (e.g., HackerOne) to operate bug bounty programs precisely to harness this third-party labor in a controlled setting. By contrast, smaller companies may have little experience with independent security research and, we hypothesized, may therefore have a tendency to act more adversarial. Our final sample includes companies with vastly different revenues; \$250 thousand to \$75 million, \$75 million to \$3 billion, and \$3 billion to \$250 billion,

³One company each from articles on NetworkWorld.com, ConsumerReports.org, LaptopMag.com, TheWirecutter.com, thoroughlyreviewed.com, top10news.com, and WirelessShack.org

Product category	Number of devices
Smart home devices	13
Multimedia devices	12
Printers	6
Tablets	6
Wearable electronic devices	6
Computers and accessories	5
Consumer software	5
Networking devices	5
Smart toys and gadgets	5
Drones	3
Video games	3
Cameras	2
Automotive GPS devices	2
Photo frames	2
Sample size	75

Table 1: The categories of target products selected for use in the study; each product is marketed by a unique contact company.

represent the approximate end points of the low, middle and top thirds of the revenue distribution.

Together, the companies selected for this study have a combined revenue amounting to over 1.5 trillion dollars, or about 8% of the U.S. gross domestic product.

2.3 Product selection

Having selected a set of target companies, we identified one product from each company that could plausibly be the target of vulnerability research. Popular and noteworthy products are particularly desirable targets for vulnerability researchers, as they might have a monetary reward attached (e.g. bug bounty programs), and could enhance the researcher's reputation. We used two distinct methods to select the products under test.

The most reviewed items on *amazon.com*. We presume the number of reviews for a product listing on *amazon.com* is a reasonable metric of product popularity. For each company identified previously, we selected the product which complied with our criteria and had the highest number of reviews, limiting the search to the first three pages after filtering by manufacturer.

Press lists of consumer products. In some limited instances, a selected manufacturer did not sell their product on Amazon, or the number of reviews did not give us confidence that the product was all that popular. In such cases, we used additional sources from the technology-consumer press, as listed above, to identify an appropriate product.

2.4 Contact selection

We hypothesized that the authorization request addressees may impact how they are handled within an organization. Requests directed to corporate counsel may engender a more legalistic response, while more informal requests of broader scope directed to the personnel responsible for software security may not. To illuminate that question, we prioritized technical managers for the contact

employee selection, and measured how many of the contacts were (explicitly) forwarded to legal. Of the responsive companies (n=30), 70% had a technical contact employee.

LinkedIn Search. LinkedIn is a professional network with over 130 million members in the United States [30]. The platform had several distinct advantages for our study:

- It established the presence of a U.S.-based workforce that could appropriately handle a request, given our objective of illuminating federal policy;
- The profiles there contain a range of self-reported, publicly-available data points about employees in our target companies, which might help us disentangle the impact of the request addressee on the request outcome;
- Finally, the network's extensive footprint make it suitable for finding people at organizations ranging from early-stage startups to very large corporations.

Our search protocol was relatively simple: we used the "advanced search" feature and filtered by current company and location (U.S. only). We prioritized contacting senior management in technical functions, and populated profiles with a higher number of connections.

We used publicly available tools such as lead-generation sites [1, 10, 20], Google Chrome extensions [4, 14, 15], and email verification tools [3, 5, 7, 9, 12, 18] to find and verify email addresses for the target employees identified on LinkedIn. Some addresses produced "risky" assessments, e.g. when a mail server flags all handles as valid (catchall). In those cases, we did the verification by sending two emails from an account belonging to a fictitious person: one to the intended address; and another to a bogus address such as `sslkgjsfg8975@example.com`. In most cases, the bogus address would explicitly bounce our email, whereas the other email would not generate a bounce warning.

General Counsel or Chief Legal Officer. Our verification protocol cannot guarantee that an email was delivered to the inbox of a real person, let alone that that person opened the communication and explicitly decided to ignore it. Since we wanted to assert with higher confidence whether or not our request was processed, we sent regular letters (via registered, private courier) to the general counsel listed on the corporate website for the companies that did not respond to our email request.

2.5 Procedure

Having identified a set of target companies and then a product and contact at each company, we were ready to contact each company. In an effort to most faithfully replicate the steps an actual researcher would take, we started with the most lightweight form of communication—email—and only escalated to postal mail if necessary. We eschewed telephonic communication due to the inherent challenges in standardizing interactive dialogues in real time and to avoid imposing undue burden on the companies under test.

We divided the study into two rounds, so that, should we encounter a problem with our methodology, we could correct it for the remainder of the sample. In the first round, we randomly assigned 10 products to the senior faculty member. We encountered no problems during the first round and proceeded to randomly

Researcher type	Targets	Email only	Email then letter	No resp.
Junior faculty	29	20	2	7
Senior faculty	9	6	2	1
U.S. ind. researcher	17	3	—	14
E.U. ind. researcher	20	3	0	17(2)
All groups	75	32	4	39

Table 2: Breakdown of the method that was effective in eliciting company responses per researcher. The U.S. researcher did not opt-in to our regular mail contact stage, and two of the 17 letters mailed by the E.U. researcher were returned to sender.

Dear *Company Representative*,

My name is *Joe Hacker* and I am a security analyst at *Hacking Hackers, Inc.*, one of the nation’s leading institutions working on cybersecurity and vulnerability analysis. Currently, my team is investigating vulnerabilities of computing devices, and we would like to include your product *Sample Product* in our tests. This email is a formal request for permission to evaluate, alter, and potentially circumvent security mechanisms (as defined by the Digital Millennium Copyright Act - U.S. Title 17 Section 1201) of the *Sample Product* for legitimate research purposes.

Should you have any questions regarding this request or the nature of our research, please contact me or browse our recent publications at www.example.com.

Sincerely,

Joe Hacker
 Senior Security Analyst
 Hacking Hackers, Inc.

Figure 1: Sample first email. The underlined sentence was included in approximately half of the requests.

assign 30 products to the junior faculty member and 20 to each of the independent researchers for the second round.

Initial contact. In each round, we created a properly instantiated version of the letter shown in Figure 1 for each contact address, with the obvious adaptations for each company, product, industry and contact person. We varied only one sentence (underlined) which changed the overall tone of the request and, we hypothesized, made one version more likely to be forwarded to the legal department or its equivalent. Each researcher sent approximately the same number of emails with and without the modifying legal sentence. We then asked each researcher to send the letters to the corresponding contact address at the company using their normal email account and include us in the BCC line and asked them to forward us any replies.

We were forced to excluded five companies at this stage of the experiment for the following reasons: the junior researcher received one response that the company had sold the product line to another

company we had already contacted, the senior researcher excluded one company from the contact list because of existing work involving the company; and three of the emails sent by U.S. independent researcher bounced and we were unable to find alternative contact addresses. Table 2 shows the number of targets assigned to each researcher and the breakdown of their responses.

Email interaction. In some cases, the initial letter sent by the researchers generated a request for more information from the company. For these, we crafted a response and asked the researcher to send it to the company. Whenever the company requested more details about the testing, we had a fixed corpus of responses to draw from. The text of the first follow-up we sent was some version of:

In general, the students are interested in exploring what kinds of network vulnerabilities might exist in *product category* and will mirror the kinds of assessments that we’ve done in the past. Regarding specific tests, expect that our students will use a combination of fuzz testing against network interfaces, reverse engineering using tools like IDA Pro, runtime taint-tracking of input buffers and so on. We may publish our findings (subject to coordinated disclosure) at a technical conference.

Please do let me know if there are any other aspects of our study I could clarify further, and which steps we could take to obtain *Company’s* authorization.

We often provided a second follow-up, usually after the firm requested information about the working process of the researcher, which described the researcher’s methodology in very high-level terms. One such example with an audio device manufacturer is below.

Currently we are assessing consumer bluetooth audio devices, including speakers and headphones. Our students will be looking for vulnerabilities that can be exploited without physical access to the devices, for example infiltrations that can be accomplished from a remote network location.

In general, we are looking for vulnerabilities that could be exploited by a malicious actor to compromise the device’s correct and safe operation, and potentially access data stored in the host cellphone or computer. We follow responsible disclosure practices, and would keep *Company* informed of our findings. Our goal is to help manufacturers address previously undiscovered vulnerabilities, so we would coordinate the disclosure (either at a technical conference or elsewhere) with you.

When a company requested a phone call, which nine of them did, we ignored the phone call request (while not explicitly denying it) and simply continued to respond via email. Of these exchanges, three companies ended up approving the request anyway. The other six stopped responding to our emails, and we terminated the exchange.

Postal followup. As noted above, we also attempted to reach each company that did not respond by email by sending a registered letter to the company’s legal department by post. The content of

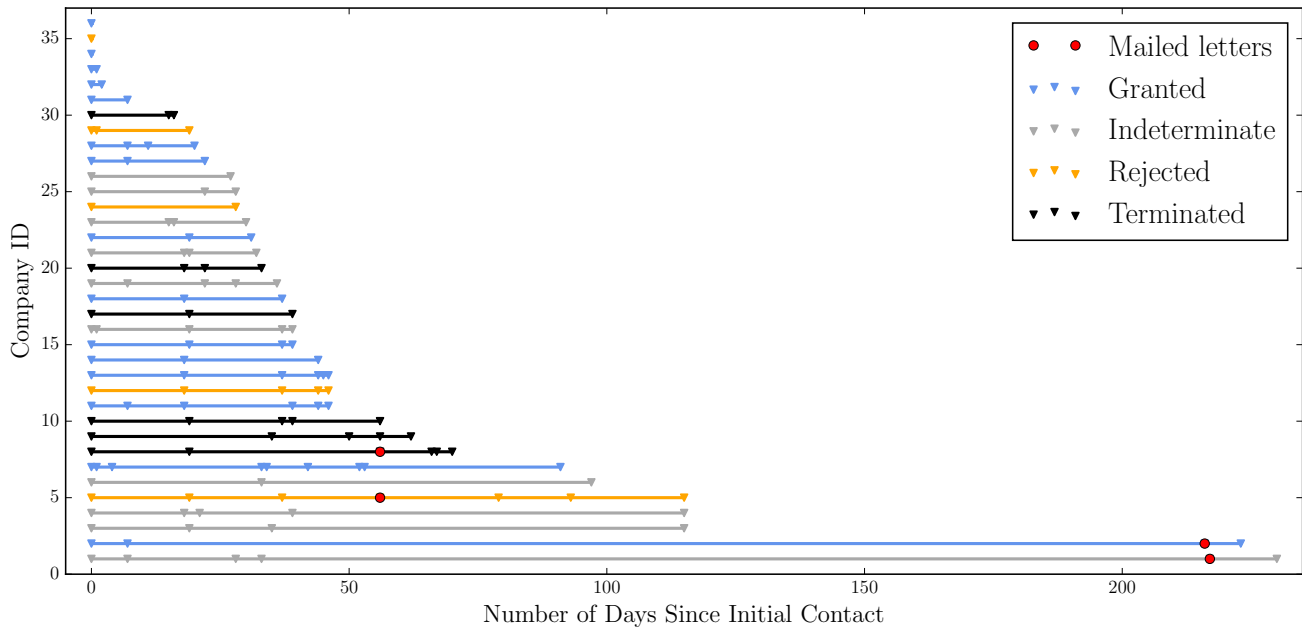


Figure 2: Timeline of researcher outgoing communications; companies are ordered by the length of the engagement. We include only the 36 companies that responded to any of our communications.

the letter was very similar to that of the email request, with a minor addition: “Please note that, despite our best efforts, we could not obtain authorization from your company by electronic means.” In 25% of cases, this resulted in an email sent by the company to the researcher. In one case, a company representative called the faculty involved and conditionally granted the request. Companies only responded to (a subset of the) letters sent by academic researchers; most were unanswered while two of the letters mailed by the E.U. researcher were returned as undeliverable.

3 AUTHORIZATION STUDY RESULTS

We now present the results of our authorization study, focusing first on the quantity and types of responses we obtained, and then analyzing whether either was influenced by our test variables (including researcher type and whether the DMCA was mentioned explicitly), as well as external factors such as company ownership and whether legal staff was explicitly engaged.

3.1 Company responsiveness

We find that, for the subset of companies that responded to our initial requests, typical response times ranged from same day to a few weeks, although several took significantly longer to resolve; e.g., one approval was not received until eight months later. Figure 2 shows this variability in response time and researcher workload. The interaction with each company is depicted as a timeline starting from initial contact; each communication is marked as a point on the line, and the entire line is colored according to the final outcome of the exchange. The four physical letters (that generated responses) are marked with red circles.

3.1.1 Classification of responses. We coded all responses to address three questions: Did the company engage with our request? If so, did the firm respond to the request definitively? And for those that granted our request, what conditions did they seek to impose on the research? These classifications result in the following seven mutually distinct result conditions:

- Unresponsive:** The company did not respond to our emails, nor postal mail to general counsel (if applicable).
- Rejected:** The company explicitly rejected our request.
- Unconditionally granted:** The company explicitly granted our request with no restrictions or requirements imposed on the research.
- Conditionally granted:** The company explicitly granted our request and imposed certain restrictions or requirements on the research.
- Indeterminate:** The company initially engaged with our request, but the exchange did not conclude with a definitive response.
- Terminated:** We terminated the exchange due to deviations from the prescribed protocol, usually following a phone call request.⁴

Unless indicated otherwise, in the following analysis we refer to requests that were either conditionally granted or unconditionally granted simply as granted. Further, we exclude the 16 unresponsive companies that only received electronic communications (the U.S. independent researcher did not send mailed letters, and two of the E.U. independent researcher’s letters were returned), as we

⁴We decided not to pursue phone calls for two reasons: we would not have fine-grained control over the exchange, as we had with written communications; and also to be mindful of the associated costs for the company involved.

have no explicit indication that they ignored our request. Finally, for the 36 responsive companies, we further tracked whether they responded to our email queries, or only after we sent a physical letter to general counsel (there is no need to differentiate based upon the method of communication used in their response, as all responses we received were electronic).

3.1.2 Response types. In this section we describe some common threads we found in the four terminal groups of the responsive companies: request granted, request explicitly rejected, exchange terminated, and indeterminate.

Granted ($n = 15$). Positive responses spanned the gamut of both enthusiasm and interest. For example, some companies responded matter-of-factly (“You are welcome to experiment on it”) or referred us to their publicly posted policies (“Since we are not looking at specific legal entanglements or tight timeframes, then the targets you’re looking for fall under our Bug Bounty and responsible disclosure policy.”) Many expressed concern that our study not impact their systems or customers (“When investigating a vulnerability, please, only ever target your own accounts and products. Never attempt to access anyone else’s data and do not engage in any activity that would be disruptive or damaging to your fellow users or to *Company*.”) Others were explicitly supportive, and some even hoped to turn our request into a sales opportunity: “If you guys are interested in buying a *targeted product* on a discounted price...”

While a minority of the companies that granted our request (13.3%) did not impose any conditions on the researcher’s testing, four fifths of the respondents explicitly requested advance warning in their communication with the researcher, with varying degrees of forcefulness, ranging from, e.g., “It’s ok with us, you can perform the tests. It will be great if you could share some of the results with us” to “...provide *Company* with written notice of any security vulnerabilities identified at least 60 days before making the work/conclusions public...” and “it’s at *Company*’s sole discretion about what...can be published.” Indeed, four companies requested editorial control over the eventual published report of the findings.

Finally, two of the companies required that the researcher sign an NDA, which often had explicit documentation requirements, such as:

Company would like to ask that you and your team: sign the attached NDA, provide *Company* full disclosure with technical details on what you were able to do and what tools/techniques were used to circumvent the controls on the tested *product*...

Rejected ($n = 5$). Companies that explicitly denied the researcher’s request disclosed very little information about their reasoning. For instance, 40% of them declined to “participate” in our study, even though no cooperation was requested by the researcher. Another 40% mentioned internal policies and constraints, to varying degrees of specificity. E.g., “We have internal constraints and contract restrictions that prevent us to give you authorization” and “We have discussed this internally and regret we must decline your request. As a policy, we don’t give permissions outside company-sponsored hacks.” Finally, one company seemed to imply they deemed our request overly broad:

As you can imagine, product security is of utmost importance to *Company*, our customers and educational partners. After careful consideration, we have made the decision to decline your blanket request to waive the rights afforded to *Company* under the Digital Millennium Copyright Act.

Indeterminate ($n = 10$). We were forced to code a significant fraction of our interactions as ‘indeterminate’ because the conversation seemed to be productive, but then dropped off and we were unable to re-engage. Common responses of this flavor include “I’ll work with the team to get you a response” and “I will forward this over to the right people to see if they are interested.” Some seemed genuinely interested, but were perhaps overruled by superiors:

We are always happy to work with security researchers. Let me know what I can do to help. I’ve included...our CTO on this, as well as...who runs our mobile group. The three of us are...responsible for security of *Company*’s products.

Terminated ($n = 6$). Most of the companies in this category (83.3%) requested a phone call, and became unresponsive when the researcher failed to arrange one. Separately, half of these companies required further clarification of the testing plans and publishing (perhaps as part of their phone call request), and were not satisfied with the written details we sent. Perhaps not surprisingly, lawyers were especially dogmatic in their requests:

Are you not willing to discuss this on the phone? I have a number of specific questions which your emails are not answering, since “coordinate” is not very specific in its meaning, and I’m not sure what you mean by deciding these things on a “case-by-case” basis. I can’t recommend to my client that it accede to your request for permission to test its product without a firmer understanding of what input the company might have on what is said in the report about the company or its products.

As with the 16 companies for which we could not confirm receipt of our request, we exclude these 10 terminated companies from the remainder of our analysis, leaving 53 companies in our data set.

3.2 Sensitivity to test parameters

We analyzed the results to identify what factors, if any, seem to influence the type or rate of response. In this section we report on both our two explicit test parameters—researcher type and DMCA mention—as well as three uncontrolled factors that appear to have an influence on the type of response. Only the first, researcher type, appears to have a statistically significant influence; the remaining relations are intriguing but have weaker significant (likely, in part, due to small sample sizes) and may be incidental ($p > 0.2$ in both binomial tests and a binary logistic regression on the other exogenous variables).

3.2.1 Researcher type. Figure 3 shows that academic researchers are significantly more likely to receive a positive response from manufacturers (13 grants out of 33 requests, versus 2 out of 20); we evaluated this assertion using a binomial test ($p = 0.028$). While it




Figure 3: Responsiveness by researcher type (n = 20, 33).

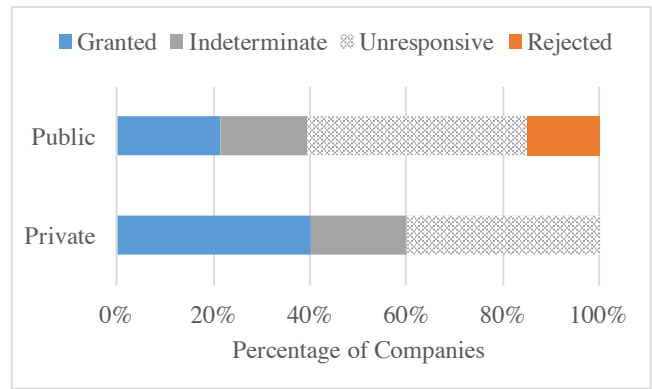


Figure 5: Responsiveness by company ownership (n = 32, 21).




Figure 4: Responsiveness by explicit mention of DMCA (n = 24, 29).

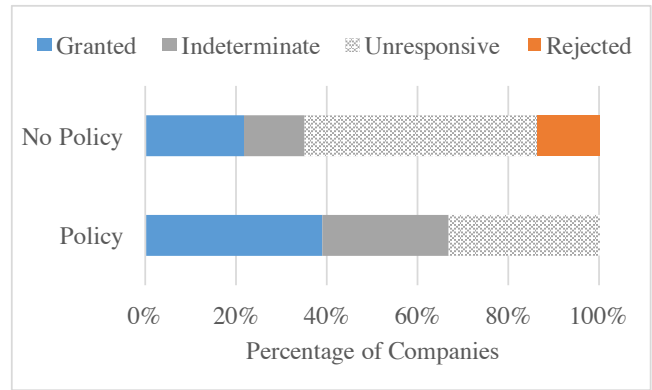


Figure 6: Responsiveness by publicly-available vulnerability disclosure policy (n = 35, 18).

is true that the only rejections we received were to requests from academics, we suspect this is biased by the fact that the academic researchers were able to engage with the companies at a higher rate; the vast majority of companies simply did not respond at all to the independent researchers despite multiple contact attempts.

3.2.2 Explicit mention of DMCA. Mentioning the DMCA explicitly did not have a meaningful effect on the likelihood that a company would initially respond to a request. Interestingly, though, the fraction of responsive companies was slightly higher for those contacts that did not mention the statute.

Figure 4 shows that explicitly mentioning the DMCA in the initial communication simultaneously reduces the likelihood that a company will grant a request (7 out of 29 versus 8 out of 24), and increases the likelihood of rejection (4 out of 29 as opposed to 1 out of 24) though, as stated previous, these differences are not statistically significant. Academic and independent researchers initially mentioned the DMCA with equal frequency, but the excluded companies resulted in some disparity (15 mentions versus 18 omissions for the academics; 9 vs. 11 for independents). The distributions compared excluded companies that were contacted by email only and did not respond, and those where we terminated the exchange.

3.2.3 Company ownership. While not an explicit test parameter in our study, we observed a difference in response rate and character depending on the type of company involved. Figure 5 shows the distribution of response classifications per company ownership. We find that the 21 private companies in our study—regardless of size—are far more likely to grant a request than the 32 public companies, although the differences are not large enough to establish statistical significance. Here again the distribution was slightly skewed between academics (24 public, 9 private) and independents (8 public, 12 private).

3.2.4 Vulnerability disclosure policy. It seems logical that companies that have already thought through the cost/benefit trade offs inherent in third-party vulnerability research are likely to be in a better position to handle our requests. Moreover, we suspect companies that provide public evidence of having conducted such a process, such as an explicit vulnerability disclosure policy or bug bounty program, are even more likely to be predisposed to consider our requests favorably. Hence, we classified companies based upon the availability of a vulnerability disclosure policy. We find that the 18 firms with such a policy are both more likely to grant our request, and less likely to reject it than the 35 companies in our data set that do not (although without a strong enough difference to establish

this finding at a significant level). Figure 6 shows the distribution of responses; academics contacted 18 companies without a policy and 15 with, independents 17 and 3, respectively.

3.2.5 *Forwards to legal departments.* Finally, while we have no way to determine which companies forwarded our request to or otherwise conferred with counsel before (or to determine whether to) responding to our request, five companies included an email chain in their response that indicated our request was explicitly forwarded to their legal team. Four of them received a contact where the DMCA was explicitly mentioned. Three of them granted the request (conditionally). One was indeterminate. One got terminated due to a phone call request. One was contacted by senior faculty via phone call (request granted). Three were contacted by junior faculty (one granted, one indeterminate, one rejected). One was contacted by the E.U. independent researcher (granted).

4 SENTIMENT SURVEY

While our authorization study indicates that a great many companies are loathe to remove legal risks from security researchers (particularly those without an academic institution backing them) this does not necessarily mean that these risks are foreclosing potential research. Indeed, measuring a “chilling effect” poses a challenge: a researcher may decide not to carry out a security evaluation for a combination of reasons whose disparate influences may be hard to untangle even for the researcher herself. We set out to measure both the perceived risk of legal action as well as the real incidence of legal threats by surveying⁵ a broad population of security researchers.

4.1 Survey methodology

To contact as many vulnerability researchers as possible, we collected the set of participants in publicly-available archives of vulnerability-related mailing lists [16] that were no older than 2014. We then inspected this set and removed any individuals, such as well-known journalists, that we knew were not themselves engaged in vulnerability research. In an attempt to further exclude individuals who do not self-identify as being engaged in vulnerability research, the first page of the survey asks:

- Q1. Does your research include the study of computer security vulnerabilities?

If the participant answers *No*, the survey ends. Otherwise, the survey proceeds to the next page. The text of the participant recruiting message was

Dear *Participant*,

We are a group of faculty and researchers at UC San Diego’s Computer Science and Engineering Department. Currently, we’re doing a study on the challenges faced by security vulnerability researchers, and would greatly appreciate your help. Please consider filling out this short survey*, to help us understand your experience. Estimated time to completion is 1–5 minutes.

⁵As with our evaluation of company responses, our survey also received IRB approval, and was declared exempt due to the low risk and anonymous nature of the survey.

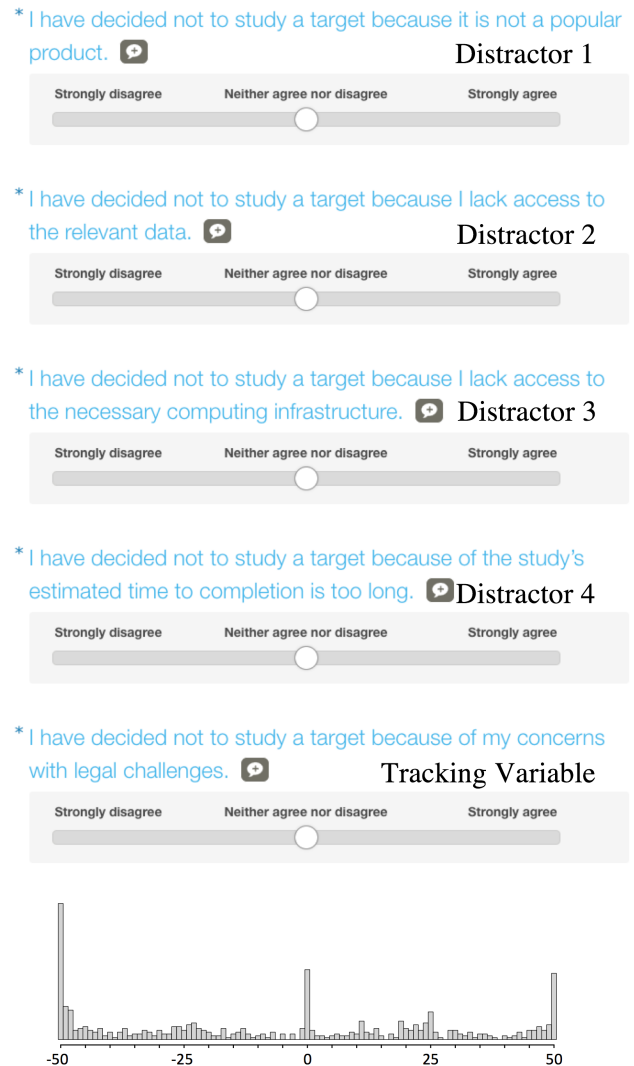


Figure 7: Top: Likert scale, which participants used to express agreement or disagreement with five statements (appearing in random order) regarding impediments to vulnerability research: four distractors and the tracking variable, “concerns with legal challenges.” Each slider had to be clicked on in order to advance to the next page, so participants had to explicitly leave the bar in the middle to continue. Bottom: histogram of the aggregate responses to all of the Likert-scale questions.

Apologies if you receive this message more than once; please ignore any additional copies of this request.

More information about our research is available at <http://www.evidencebasedsecurity.org/>

where the footnote was a hyperlink to Question 1 (collected via the SurveyMonkey platform).

4.1.1 *Relative influence of legal risks.* To determine whether the perceived risk of legal action was a deterrent, we asked survey participants a series of five questions about factors that might influence their decision to undertake a particular vulnerability research project. Both the questions (which appeared in randomized order) and the response distributions are shown in Figure 7; answers were recorded on a Likert scale [11] with extremes marked *Strongly disagree* and *Strongly agree* and a neutral position marked *Neither agree nor disagree*. While potentially interesting in themselves, for our purposes four of these five questions serve simply as distractors that allow us to calibrate the responses to the tracking question (“concerns with legal challenges”) and to avoid leading the participants.

4.1.2 *Past fear of legal risks.* After responding to the first set of Likert-scale questions, participants are taken to a new page and asked two sets of yes/no questions, in order to gauge both their subjective evaluation of the impact of legal challenges on their research, and their factual experience with legal threats or actions. In the first set, participants are asked:

- Q7a. Have you ever feared legal action (e.g. a cease-and-desist letter, or a civil lawsuit) as a consequence of your vulnerability-related research?
- Q7b. Has this fear of legal action prevented you from engaging in, or prompted you to modify, a research project? (shown if answer to Q7a was *Yes*)

Participants that answered *Yes* to the second question were offered a chance to explain their answer in a text box.

4.1.3 *Experiences of legal threats and action.* The second set of yes/no questions attempted to elicit factual information about legal action faced by the participant:

- Q8a. Have you ever been threatened with legal action (e.g. have you received a cease-and-desist letter from a manufacturer or copyright owner) as a consequence of your vulnerability-related research?
- Q8b. Have you ever been a named defendant in a court of law (e.g. have you been sued) as a consequence of your vulnerability-related research?

Participants that answered *Yes* to either question were offered a chance to describe their experience in a text box.

4.2 Responses

Of the 1,369 individuals invited to participate in the survey, 158 (11.5%) visited the survey page and 139 (88%) of those responded *Yes* to the first, qualifying question self-reporting as vulnerability researchers. Some of these respondents did not complete any further questions and were removed from the sample, leading to a final set of 110 respondents.

4.2.1 *Relative influence of legal risks.* In order to evaluate whether or not researchers were likely to single out legal challenges as an obstacle to research, we compare the distribution of the tracking variable with the combined distribution of the four distractors. The continuous slider provided integer response values in the range [-50, 50] as shown on bottom of Figure 7. The responses

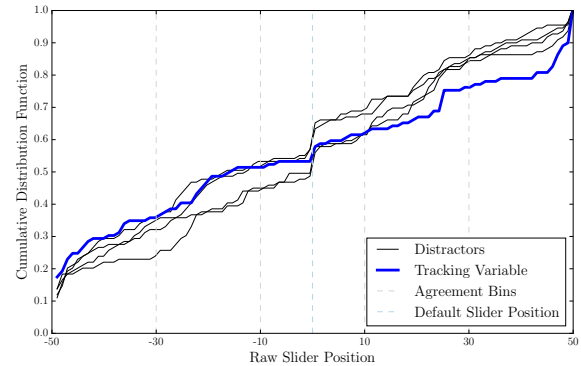


Figure 8: The empirical cumulative distribution functions of the raw Likert-scale observations for the tracking variable and each of the distractors. Dashed vertical lines depict the bin cutoffs as well as the starting slider position.

show clear modes at the extremes (-50 and 50) and at 0, with less pronounced modes at values between the extremes and 0. Figure 8 shows the cumulative distribution of Likert-scale responses for each of the four distractors (black lines) and the tracking variable (blue line), which further illustrates the modes. The presence of these modes strongly suggest the observations can be binned into five discrete categories corresponding to the extremes, 0, and values between the extremes and 0. Figure 9 shows the responses binned into these five categories with breaks at -30, -10, 10 and 30.

The top four histograms in Figure 9 show the distribution of the distractor variables and the fifth the distribution of the tracking variable. The bottom histogram shows the difference in the response counts between the tracking variable and the average response count for the same bin across all four distractors. For example, there were an average of 35.5 responses that fell into the *Strongly disagree* bin received across the four distractor variables, and 40 *Strongly disagree* responses for the tracking variable; the difference, 4.5, is shown in the first bar. Figure 9 also shows the 95% confidence interval for the responses as lines extended above and below the top of the histogram. For any given response bin, the confidence interval shows the range of possible response frequencies for which the observed mean has a probability of at least 5%.

Participants were 1.5× more likely to indicate strong agreement with the statement they chose not to study a target because of concerns with legal challenges than they were for the distractors, a statistically significant difference ($p = 0.034$). In fact, for 14% of participants the tracking question was the only *Strongly agree* among the five Likert-scale questions.

We also observe that the tracking variable produces a more polarized response, with a greater number of responses falling into the *Strongly disagree* or *Strongly agree* bins than for the tracking variables. In order to statistically test that increased polarization in the raw distributions, we used a Levene’s test (carried out on the full [-50, 50] range, rather than binned responses) which showed a statistically significant ($p < 0.001$) difference in the variance of the tracking variable compared to the distractor variables




Figure 9: The top five histograms show the distribution of binned responses for the four distractors (in blue) and the tracking variable (in orange). The bottom chart plots the mean difference in responses between the tracking variable and distractors for each bin; whiskers show the 95% confidence interval.

treated as a single distribution. Table 3 summarizes Levene’s test results. The first row of the table shows the result of comparing the tracking variable distribution T to the distribution combining all four tracking variables, and the remaining rows show individual comparisons. (The Kolmogorov-Smirnov test did not show a

Null hypothesis	Levene	
	coeff.	p-val
$T \sim D_{\{1,2,3,4\}}$	38.201	<0.001
$T \sim D_1$	3.811	0.052
$T \sim D_2$	6.999	0.009
$T \sim D_3$	4.596	0.033
$T \sim D_4$	4.159	0.043
$D_1 \sim D_2$	0.501	0.480
$D_1 \sim D_3$	0.015	0.903
$D_1 \sim D_4$	0.007	0.932
$D_2 \sim D_3$	0.380	0.538
$D_2 \sim D_4$	0.389	0.534
$D_4 \sim D_4$	0.001	0.974

Table 3: Statistical analysis of continuous-slider Likert-scale responses.

statistically-significant difference between the distribution of the tracking variable and the combined distribution of the distractor variables; we suspect that the difference in the CDF seen in Figure 8 for the tracking variable could not be captured by the test, as the tracking distribution departs from the distractors for a constrained set of values at the extremes.)

4.2.2 Past fear of legal risks. When asked if they had feared legal action as a consequence of their vulnerability-related studies (Q7a), about half of respondents answered *Yes* (49.1%) and half *No*. This question differs from the tracking question in an important way: (Q7a) asks if the researchers feared legal action, while the tracking question on the Likert-scale survey asked whether any such concerns actually *caused* the respondent not to study a target. Of those who answered (Q7a) in the positive, the breakdown on the tracking question was: 25.4% *Strongly disagree*, 5.5% *Disagree*, 14.5% *neutral*, 20.0% *Agree* and 34.5% *Strongly agree*. That is, 30.9% decided to study a target despite fearing legal action. Among those who indicated they did not fear legal action, a quarter either agreed or strongly agreed that concerns about legal action had made them decide not to study a target.

If a respondent answered (Q7a) in the positive, we followed up with (Q7b), asking whether fear of legal action caused him/her to modify a project. 52.7% answered *Yes* and 47.3% answered *No*. This question more closely aligns with the Likert-scale tracking question in meaning. Indeed, 55.5% who agreed or strongly agreed with the tracking question answered *Yes* to (Q7b). Conversely, 69% of those who answered *Yes* to (Q7b) either agreed or strongly agreed with the tracking question.

4.2.3 Experiences of legal threats and action. When asked if they had been threatened with legal action (Q8a), 22% of respondents answered *Yes*. Of those who answered (Q8a) in the positive, the breakdown on the tracking question was: 29.2% *Strongly disagree*, 4.2% *Disagree*, 12.5% *neutral*, 4.2% *Agree* and 50% *Strongly agree*. Indeed, the median response to the tracking question was *Strongly agree* for this group of respondents.

Two respondents answered *Yes* to (Q8b), stating that they had been a defendant in a court of law. One of them responded *Strongly agree* to the tracking question, and the other one was neutral. Naturally, both of these participants also answered *Yes* to (Q7b), stating that legal concerns had factored in their engagement with a vulnerability-related research project.

4.2.4 Qualitative experiences of legal threats. Those participants who either indicated that they had chosen not to proceed with their research due to fear of legal action or those who had experienced concrete legal actions were further asked to provide free form text describing how this impacted them.

For those who reported that they had modified their research due to fear (over 52%) the most common trait was the choice to discontinue research on such targets (over 15 participants). In some cases these reflect the classic notion of “chilling effects” as they apply to potential future research (“I avoid researching vulnerabilities in those products owned by companies who threaten security researchers with legal action”, “Identifying whether legal consequences are likely is a step I have planned when evaluating potential avenues for research”, “always get written consent before doing anything”), but in other cases this sentiment was applied retroactively and has kept completed or in-progress research from being shared, even with the vendor (“We removed the relevant research files from our group site”, “[I] destroyed my notes, and did something else”, “I sat on a vuln for a decade”). Finally, in several cases the researcher did release their information in spite of fears of legal risks, but reported trying to protect themselves by doing so anonymously; for example reporting that they “released info anonymously”, “gave my data anonymously to a journalist”, or discussing their “anonymously-sourced exploit”.

Those who experienced legal action (22%) ranged from “e-mails with veiled legal threats, citing violation of licensing terms and conditions”, to concrete threats of lawsuits under particular statutory provisions (“contacted by the manufacturer’s legal team and threatened with multiple lawsuits... [including threats of] arrest for violating 18 USC 1030”). The comments also illustrate the range of complexities encountered given that companies, researchers and researcher’s employers may be in different countries operating under different legal regimes and multiple researchers reported receiving legal threats from overseas companies (and at least one viewed their location as a form of defense, “you are welcome to sue a private person in Russia”). Another researcher identified the mixed messages afforded to some vulnerability researchers (“[received] warnings that I would be arrested if entering China. Followed an invite by the company to come to China and speak.”).

Responses that discussed outcomes were split between admissions of capitulation (“had to drop the name of the product”, “removed articles/files”, or “led to non-action by the vendor and made us not disclose the vulns”) and resistance (“was meant to scare me into shutting up. So, I obviously did the exact opposite” or “I told them to f[...] off”). Multiple participants mentioned communicating with their own lawyers or EFF (the only external organization mentioned) to obtain legal advice and one mentioned deescalating the legal conflict through a person that they knew personally at the company. Just under 2% of survey participants reported being a named defendant in a court of law.

5 DISCUSSION

The ultimate goal of any study as ours is to provide data to frame further policy discussion—in this case on the question of whether legal threats significantly deter third-party security research (and the related question, which we do not consider, of whether this problem should be addressed). In particular, we have been motivated to help resolve a common rebuttal exemplified by the comments of the Business Software Alliance (BSA) in the most recent DMCA rulemaking concerning a security research exemption [25]. In their comments, the BSA argues that “The proponents of the exemption have put very little in the record to demonstrate that researchers are currently declining to engage in good faith security testing as a result of the prohibition against circumvention of access controls contained in [the DMCA]” and further that “Almost all software companies already have in place carefully tailored processes for identifying vulnerabilities and working with independent researchers and members of the public to address them”. In our work we have provided an empirical basis for evaluating both questions.

To the latter, it is certainly true that a broad array of companies have concrete vulnerability disclosure policies (if not explicit remuneration via bug bounty programs) that can implicitly encode safe harbors for researchers who follow their disclosure requirements. At the same time, our data indicates that this is far from a pervasive legal posture and only a minority of companies contacted were willing and able to present such terms when explicitly asked for permission. Moreover, there is significant skew in the response as a function of the requester with academic researchers being three to five times more likely to receive a response than independent researchers and six times more likely to receive a positive response (i.e., a grant of permission).

The reasons behind this discrepancy are not completely clear, but we suspect it may reflect a number of factors including the reduced personal risk in the university setting (i.e., that university researchers are less likely to be intimidated by legal threats given that their institution’s general counsel will defend them), the worse “optics” associated with pursuing legal threats against educational institutions vs. individuals, the desire to recruit from university campuses and sentiments that academics may be more amenable to coordinated disclosure requirements. The consequence is that there is implicitly a system of legal haves and have-nots, with academic researchers being doubly favored, receiving better treatment by companies and being better shielded from personal liability.

To address the issue of whether these factors engender chilling effects our research shows that while legal risks do not drive the decisions for a majority of researchers surveyed, such concerns are not insignificant. Indeed, a substantial minority (over 38%) agreed or strongly agreed with the statement that concerns with legal challenges had caused them to not study a particular target and almost a quarter (22%) had experienced legal threats as a direct result of their vulnerability research work.

Taken together, we believe our data provides an empirical basis to argue that legal risks (both perception and reality) remain an issue of concern for the security research community. Moreover, these risks may disproportionately affect the experience of non-academics, a population of researchers that may be less well represented in the circles seeking to influence public policy.

6 RELATED WORK

There is considerable discussion of the “chilling effects” of U.S. and E.U. laws on industry activity, notably describing unforeseen consequences of regulation [23, 29]. Others have quantified DMCA copyright take down notices [22]. However, we are unaware of any empirical studies of security vulnerability disclosure.

The value of independent third-party vulnerability research has been widely discussed in the literature. For example, Rand recently released a study on zero-day vulnerabilities, where they found that their average life expectancy is high, at 6.9 years. The authors conclude that “Defenders... likely will want to disclose and patch a vulnerability upon discovery” [21], highlighting the critical nature of security vulnerability research. Others have argued that vulnerability discovery may not be beneficial (i.e., if software defect rates remain high) [33] and questions of these processes (vulnerability density and rediscovery probability) are of considerable interest to both the research and policy communities in this space.

Researchers have also studied vulnerability reward (bug bounty) programs, particularly their cost effectiveness [28]; these by definition do not include vendors without a vulnerability rewards program. Although consumer-facing companies have certainly created such programs in recent years, they are far from universal.

Finally, a related issue is the research community’s (often incomplete) understanding of the regulations: “uncertainty among researchers about what the law actually says, as well as doubt about the ethics of some activities, may hold back certain research efforts” [26]. Further, Ohm finds that while recent network measurement studies have taken steps towards legal compliance, “many of these papers may fall short of legal expectations.” Most pressingly, the authors argue that “compliance with the law might constrain measurement methodologies in ways completely at odds with the goals of most research; and second, while many of the rules will [limit exposure], few will reduce the risk of liability to zero” [32].

7 CONCLUSION

In this work, we took a first step toward characterizing the chilling effect of legal threats to vulnerability researchers. To do so, we carried out an empirical study of companies’ willingness to grant security researchers permission to conduct a security evaluation of their products. We found that researcher affiliation played a major role in the success of such a request: about 40% of companies asked by researchers with an academic affiliation granted permission, while less than 10% of companies grant permission to independent researchers. We also conducted a survey of security researchers and found that nearly a quarter strongly agreed with the statement that they had decided not to study a target because of concerns over legal challenges. Other factors, such as time and resource constraints, played a less important role. Finally, and surprisingly to us, we also found that 22% of researchers reported receiving legal threats because of their vulnerability research.

ACKNOWLEDGMENTS

This work is supported by grants from the National Science Foundation (CNS-1237264) and the William and Flora Hewlett Foundation (2016-3838). We are indebted to the anonymous academic and independent researchers that agreed to contact the companies as part of our vulnerability study, as well as the members of the community

that responded to our sentiment survey. We also thank the CCS reviewers for their comments on an earlier draft of this manuscript.

REFERENCES

- [1] Data.com connect - a Salesforce product. Available online at <http://connect.data.com>. Accessed April 2017.
- [2] Dow Jones U.S. technology index - S&P Dow Jones indices. Available online at <http://us.spindices.com/indices/equity/dow-jones-us-technology-index>. Accessed April 2017.
- [3] Email checker. Available online at <http://email-checker.net>. Accessed April 2017.
- [4] Email hunter. Available online at <https://hunter.io/chrome>. Accessed April 2017.
- [5] Email verifier. Available online at <https://hunter.io/email-verifier>. Accessed April 2017.
- [6] Fortune 500. Available online at <http://fortune.com/fortune500/>. Accessed June 2016.
- [7] Free email verifier. Available online at <http://verify-email.org>. Accessed April 2017.
- [8] iShares U.S. technology ETF | IYW. Available online at <https://www.ishares.com/us/products/239522/ishares-us-technology-etf>. Accessed April 2017.
- [9] kickbox. Available online at <https://kickbox.io>. Accessed April 2017.
- [10] lead411. Available online at <https://lead411.com>. Accessed April 2017.
- [11] The Likert scale explained. Available online at <https://www.surveymonkey.com/mp/likert-scale/>. Accessed November 2015.
- [12] Mail tester. Available online at <http://mailtester.com>. Accessed April 2017.
- [13] A qualified win for cybersecurity researchers in DMCA triennial rulemaking. Available at <https://cdt.org/blog/a-qualified-win-for-cybersecurity-researchers-in-dmca-triennial-rulemaking/>. Accessed April 2017.
- [14] Reachout. Available online at <http://www.zoominfo.com/business/reachout-info>. Accessed April 2017.
- [15] Rocketreach. Available online at <http://rocketreach.co>. Accessed April 2017.
- [16] Seclists.org security mailing list archive. Available online at <http://seclists.org>. Accessed November 2015.
- [17] Unintended Consequences: Fifteen Years under the DMCA. Available at <https://www.eff.org/pages/unintended-consequences-fifteen-years-under-dmca>. Accessed April 2017.
- [18] Verify email address. Available online at <http://www.verifyemailaddress.org>. Accessed April 2017.
- [19] Y combinator. Available online at <http://www.ycombinator.com/>. Accessed April 2017.
- [20] zoominfo. Available online at <http://www.zoominfo.com>. Accessed April 2017.
- [21] Lillian Ablon and Andy Bogart. *Zero Days, Thousands of Nights*. RAND Corporation, Santa Monica, CA, 2017.
- [22] Aleecia M. Aleecia M. McDonald and Wendy Seltzer. Quantifying the Internet’s Erasers: Analysis through Chilling Effects Data. *Privacy Law Scholars Conference*, 2013.
- [23] Barendt, Lustgarten, Norrie, and Stephenson. *Libel law and the media: the chilling effect*. Clarendon Press, 1997.
- [24] BlackRock. About us. Available online at <https://www.blackrock.com/corporate/en-us/about-us>. Accessed April 2017.
- [25] BSA. Long Comment Regarding a Proposed Exemption Under 17 U.S.C. 1201. 2014.
- [26] Aaron J. Burstein. Amending the ECPA to enable a Culture of Cybersecurity Research. *Harvard Journal of Law & Technology*, 22(1):168–221, 2008.
- [27] Center for Democracy & Technology. The current DMCA exemption process is a computer security vulnerability. Available at <https://cdt.org/blog/the-current-dmca-exemption-process-is-a-computer-security-vulnerability/>. Accessed April 2017.
- [28] Matthew Finifter, Devdatta Akhawe, and David Wagner. An empirical study of vulnerability rewards programs. In *Proceedings of the 22nd USENIX Security Symposium*, pages 273–288, 2013.
- [29] Thomas W. Hazlett and David W. Sosa. Was the Fairness Doctrine a “Chilling Effect”? Evidence from the Postderegulation Radio Market. *The Journal of Legal Studies*, 26(1):279–301, 1997.
- [30] LinkedIn. About us. Available online at <https://press.linkedin.com/about-linkedin>. Accessed April 2017.
- [31] John Markoff. Record panel threatens researcher. *New York Times*, April 2001.
- [32] Paul Ohm, Douglas Sicker, and Dirk Grunwald. Legal Issues Surrounding Monitoring During Network Research (Invited Paper). *7th ACM SIGCOMM conference on Internet measurement*, 2007.
- [33] Eric Rescorla. Is finding security holes a good idea? *IEEE Security and Privacy*, 3(1):14–19, January 2005.
- [34] U.S. Copyright Office Library of Congress. Exemption to Prohibition on Circumvention of Copyright Protection Systems for Access Control Technologies, 2006. Available at <https://www.copyright.gov/fedreg/2006/71fr68472.html>. Accessed April 2017.