"*Sticking their heads out above the parapets*": Lived Experiences of Legal Risks in Research (Extended)

Sunoo Park
New York University

Daniel R. Thomas
University of Strathclyde

Full Version of USENIX Security 2026 Publication[1]
See our Lawfare post for a short-form overview.[2]

Abstract

Overbroad computer crime, intellectual property, and other laws are well known to create legal risks that can discourage essential research. Notable examples include the US Computer Fraud and Abuse Act and the UK Computer Misuse Act. Because such laws fail to distinguish malicious hacking from good-faith testing and research, researchers face serious legal risks for public-interest research activity like identifying software or hardware vulnerabilities or scraping data. Despite the research community's broad awareness of these risks, our understanding of their practical impacts is limited, as most of the community's knowledge comes from anecdotal evidence rather than systematic study.

We conduct the first qualitative study focused on researchers' lived experiences, to empirically document the *impacts of legal risks and threats* on research and researchers, and *how researchers navigate legal risk situations*. Our study engages two participant groups: researchers with legal-risk experiences in the UK or the US ($N_R = 36$), who discuss 130 projects and incidents spanning over three decades, and professionals that offer support to researchers navigating legal risks ($N_S = 8$), who have collectively supported thousands of researchers. We thus provide an unprecedented big-picture view of researchers' experiences with legal risks. We synthesise actionable strategies for researchers, and our findings provide evidence to support policy reform.

Contents

1 Introduction **4**
2 Methodology **5**
3 Analysis **9**
3.1 The Big Picture . . . 9
3.1.1 Legal Risks Are All Around Us . . . 9
3.1.2 Research Areas & Trends Over Time . . . 10
3.1.3 Researchers' Awareness & Motivations . . . 13
3.1.4 Legal Threats & Legal Representation . . . 14
3.2 Impacts on Research(ers) . . . 15
3.2.1 Feelings . . . 15
3.2.2 Chilling Effects . . . 16
3.2.3 Stockpiling Vulnerabilities . . . 17
3.2.4 Personal and Health Impacts . . . 18
3.2.5 Professional Impacts Beyond Chilling Effects . . . 18
3.2.6 Diversity of Researcher Risk Profiles . . . 19
3.3 Strategies, Precautions, Advice . . . 21

[1] https://www.usenix.org/system/files/usenixsecurity26-park.pdf
[2] https://www.lawfaremedia.org/article/the-legal-risks-that-chill-good-faith-security-research

3.3.1 Advice for Researchers . . . 24
3.4 Lawyering Up . . . 26
3.4.1 Experiences With Lawyers . . . 27
3.4.2 Institutional Lawyers . . . 28
3.4.3 Financial and Resource Constraints . . . 30
3.4.4 More From Supporting Professionals . . . 31
3.5 Institutional Support . . . 32
3.5.1 Institutional Support in Industry . . . 32
3.5.2 Institutional Support in Academia . . . 32
3.5.3 Note on Independent Researchers . . . 34
3.6 Jurisdictional Variation . . . 34
3.6.1 Legal Threats & Legal Representation . . . 34
3.6.2 Types of Research . . . 34
3.6.3 US vs. UK Computer Crime Law . . . 34
3.6.4 Federal vs. State Within The US . . . 35
3.6.5 Crossing Jurisdictions . . . 35
3.7 Notable Areas of Research . . . 35
3.7.1 Vulnerability Research . . . 36
3.7.2 Social Computing . . . 44
3.7.3 Network Analysis . . . 45
3.7.4 Cryptography and Export Controls . . . 45
3.7.5 Malware Writing . . . 45
3.7.6 Interview Research . . . 46
3.7.7 Incident Response . . . 46
3.8 Ethics Review . . . 46
3.8.1 Institutional Review . . . 46
3.8.2 Ethics Review at Conferences . . . 47
3.9 Other Notable Themes . . . 49
3.9.1 Academic Freedom . . . 49
3.9.2 Non-Disclosure Agreements (NDAs) . . . 49
3.9.3 Advocacy for Change . . . 50
3.9.4 UK Trusted Research . . . 50
3.9.5 Non-Legal Risks . . . 51
3.9.6 Other Concerning Experiences . . . 52
4 Recommendations **52**
4.1 Researchers . . . 52
4.2 Other Stakeholders . . . 53
5 Discussion **54**
6 Related Work **55**
7 Conclusion **56**
Acknowledgements **56**
A Useful Resources for Researchers Facing Risk **56**
B Additional Ethical Considerations **57**
C Glossary of Acronyms **58**
D Index of Referenced Quotes **58**
E More Details On Participants and Methodology **59**
F Example: An Ounce of Prevention **60**
G Background on UK Ethics Committees **61**
H Additional Plots **61**
References **64**

1 Introduction

> *« You can't really have a functional security research career without taking some level of risk; there's just not enough legal certainty. »* Q1

Anti-hacking laws aim to serve the public interest by introducing legal penalties for malicious hacking. Paradoxically, however, their ambiguous and overbroad scope can *harm* the public interest in practice. Because such laws often fail to distinguish clearly between malicious hacking and good-faith testing and research, many researchers face serious legal risks for essential research activities. As such,

> "*unfortunately, the legal system gives people that don't like your research tools to disrupt your research.*"

This legal uncertainty, in turn, creates a "chilling effect"[3] across the scientific community [4, 5]. Academics, institutions, and independent researchers may abandon critical projects that could improve our understanding of technologies that millions of people rely on every day, and could improve these technologies' resilience to errors and attacks.

Prior research and advocacy has documented the legal risk landscape (e.g., [6, 7, 8, 9, 10, 11, 12, 13, 14]), but offers limited insight on the practical impacts of legal risks and researchers' experiences of legal risk situations. Despite the research community's broad awareness of these risks, empirical evidence is scarce; indeed, there is a natural incentive not to speak out about legal risk experiences as doing so could heighten risk exposure, and much of our community awareness comes from anecdotal evidence (e.g., [15, 16]) and professional folklore rather than systematic study. Yet empirical evidence can be essential for reform [17].

Our qualitative study, focused on legal risk experiences of researchers in the US and UK, provides **the most comprehensive account to date of the serious impacts of legal risks on research and researchers in computer science**.

We interviewed $N_R = 36$ researchers with research-related legal risk experiences as well as $N_S = 8$ professionals who support researchers affected by legal risks, such as lawyers. The researchers provide essential firsthand perspectives across 130 distinct incidents; the supporting professionals offer valuable complementary perspectives, as they have exposure to hundreds of researchers' stories, and often a broader knowledge of the risk landscape than individual researchers.

We thus provide an unprecedented **big-picture view of legal risks** in our field (§3.1), including broad trends such as improvements in the legal risk landscape for vulnerability research and escalation of legal risk for social computing researchers in a changing political climate. Our analysis covers:

1. **Impacts** (§3.2), including chilling effects, researchers' feelings, personal and health impacts, other professional impacts, and a discussion of the diversity of researcher risk profiles depending on individual circumstances.
2. **Strategies and precautions** (§3.3) employed by researchers, such as around anonymity, data, contracts, legal advice, public relations, and research design.
3. **Experiences with lawyers** (§3.4), including personal and institutional lawyers, benefits and costs of legal representation, (lack of) institutional support for researchers, attitudes towards lawyers, and lawyers' perspectives.
4. **Institutional support** (§3.5) for researchers in academia, industry, and government.
5. **Jurisdictional variation** (§3.6) between the UK and US, including trends in civil and criminal legal threats.
6. **Areas of research** (§3.7), where we found vulnerability and social computing were the biggest areas by far, but our participants discussed more than a dozen research areas.
7. **Ethics review** (§3.8) processes' implications for legally risky research (including both institutional and conference review).

Our participants shared diverse experiences ranging from high school tinkering to senior professional roles, from losing their jobs to prison time, from battles in court to sparring in the press, and more. We are grateful for their candid contributions, without which this work would not have been possible.

We find the legal risk landscape has improved in many respects over the last decades. Despite this, we heard from a majority of participants that legal risk is very common or unavoidable for many researchers, that legal risk is "*always in the back of [researchers'] mind[s]*", and that "*there is still a lot of room to go, and hopefully in the future we can reform these laws.*"

The chilling effects are real (§3.2.2). Researchers recounted being driven away from legally risky research areas, and suppressing research results due to legal risk — results that may be essential to understanding, securing, and holding accountable systems that millions of people rely on.

[3] *Chilling effects* refers to a phenomenon where legal or social infrastructure "chills" or deters people from speaking or acting freely, due to the fear or risk of adverse consequences [1, 2, 3]. In our context, we focus on "chilling" of research activities: that is, people refraining from freely conducting and speaking about good-faith research activities.

Many researchers persevere despite the risks, often motivated by "*doing the right thing*" and conceptualising research as an essential public service.

> *« Researchers publish their results; bad actors don't. Because they publish their results, researchers effectively stick their heads out above the parapets in terms of legal liability. »*

By thus exposing themselves to risk to do public-interest work, researchers face potentially serious personal and professional consequences (§3.2).

> "*I sent some emails warning people of a security incident and [law enforcement] c[a]me up at me with machine guns — that kind of doesn't seem to balance.*"[4]

Vulnerability research stands out as the area with the most extensive history of legal risks. Social computing research stands out as the one area where we heard legal threats have been increasing and worsening in recent years, correlated with a rise in politically motivated threats to research over the last decade (§3.1.2).

> "*[W]e're in an environment now that not only is hyper politicised, but the biggest owners of the biggest platforms themselves have an incentive to cozy up to this administration and as a result there's sort of been . . . a full-frontal assault on those who work on platform accountability and social media issues.*"Q2

We also investigated researchers' experiences with lawyers and institutions, to whom they may turn for support. We found institutional support for researchers is highly variable, especially in academia (§3.5). Relatedly, we heard that researchers having their own lawyer can be important to navigate legal risk. Despite some initial scepticism, researchers who hired their own lawyers reported overwhelmingly positive experiences, whereas experiences with institutional lawyers were mixed (§3.4).

Jurisdictional variation was not as pronounced as one might expect for two jurisdictions with distinct laws, partly because UK-based researchers not uncommonly have to navigate legal-risk situations under US law. Notably, the absence of a civil cause of action in UK computer crime law reduces researchers' risk exposure from private parties (§3.6).

Finally, ethics review processes emerged as a recurring theme that can raise particular challenges for legally risky research (§3.8).

[4] This quote was part of a participant's literal description of a particular instance where law enforcement approached them unexpectedly.

Overall, the legal-risk experiences we document are highly individual and context-specific. Nonetheless, consolidating our findings with strategies, precautions, and advice from our participants, we offer three broad recommendations that are unilaterally actionable by researchers, suggest how a range of other stakeholders can help mitigate risks (§4), and provide discussion of broader societal context, including national security risks, the role of recent advancements in AI-powered cybersecurity tools, and advocacy and reform to date (§5).

Roadmap. §2 describes our **methodology**, §3 presents **analysis**, drawing on which §4 offers succinct **recommendations**. Then, §5 provides **discussion** and §6 covers **related work**.

2 Methodology

We followed Mason's [18] and Saldaña's [19] qualitative research methods. We considered the ontological and epistemological perspectives we wished to investigate, starting from the broad area of *researchers' legal risk experiences in vulnerability and social computing research*, based on which we developed 10 top-level research questions and 87 sub-research questions. The six questions below give an abbreviated summary; full details are in [20].

RQ-A. How do legal risks impact researchers' ***research*** topics and processes?

RQ-B. How do legal risks impact ***publication*** and ***communication about legal risks***?

RQ-C. What is the role of researchers' ***institutions***, ***lawyers***, and other parties?

RQ-D. What relevant ***community norms & practices*** exist?

RQ-E. Are there ***patterns*** in legally risky research projects that go well or badly?

RQ-F. How do recipients ***react*** to researchers' disclosures, and how do researchers ***navigate (difficult) disclosures***?

From our research questions, we developed two sets of interview questions (given in [20]), one for researchers, and one for people who support researchers (respectively, 12 and 8 high-level questions with sub-questions). We used these questions to conduct semi-structured interviews by video/audio call or in person, which took place between November 2024 and January 2026 and lasted between 30–120 minutes (on average, 75 minutes), with 54 hours of interviews in total. Some interviews were conducted in two sessions. Both authors were at all but 3 interviews, which were run by just one author.

Recruitment. Potential participants were identified through public documents that indicated they

Experience	≤5	6–10	11–15	16–20	>20
Count	3	6	10	5	10

Table 1: Researchers' years of relevant experience

Experience	≤10	>10
Count	1	7

Table 2: Supporting professionals' years of relevant experience

Jurisdiction	UK	US	Both
Count (Researchers)	12	22	1
Count (Supporting Prof'ls)	6	0	2
Count (Projects)	36	84	–

Table 3: Participants and projects by jurisdiction. Entries under "Both" are *not* also listed under "UK" and "US". Projects were labeled with their primary jurisdiction, never as "Both".

	Number of Quotes Per Researcher	Supporting Prof'l
Minimum	1	3
Mean	7.0	10.9
Median	4.5	9
Std. Dev.	5.2	7.8

Table 4: Participant quotes in this paper

had relevant experiences, through the authors' professional networks, through snowball sampling, and through outreach to relevant researcher communities (e.g., via online forums and conferences).

Participants. We interviewed $N_R = 36$ *researchers* with US- and/or UK-based research-related legal risk experiences as well as $N_S = 8$ people who support or have supported researchers affected by legal risks (whom we call *supporting professionals*), such as lawyers; who have directly supported hundreds of researchers. In total, $N = 44$.

22 of our researcher participants were based in academia; some of these also had industry experience. We additionally had 14 independent or industry researcher participants. Three participants also had experience with government research.

Researchers' self-disclosed genders were 17 male, 7 female, and 12 did not disclose or said they would prefer not to.[5] Recent estimates suggest women make up around 15–30% of computer-science researchers [21] and 11–24% of the cybersecurity workforce [22], so our numbers roughly track the gender imbalance in the field. For anonymity, we use *they* pronouns for all participants except in one quote (Q14) where the participant's gender provides important context.

Participants discussed experiences ranging from high school to independent research to senior positions in industry, academia, and government. Table 1 summarises their length of experience. All supporting professionals had directly helped at least 20 individual researchers. Some directly helped hundreds; some impacted thousands.

Scope of Research. Anyone who self-identified as a researcher in security or measurement was treated as eligible. This included people who engaged in research in freelance or hobbyist capacities as well as professionals. The result is a broad construal of "research" encompassing a wide range of activities experimenting with or investigating how software and hardware systems, products, and services work.

Quotes. Participant quotes are styled in *italics*. To protect anonymity, quotes are not linked to individual participants — not even by pseudonym, as this raised concerns of reidentification by associating quotes from the same participant — and every quote has been explicitly approved by the source participant. Quotes from supporting professionals appear « *in blue with guillemets* »; all other quotes are from researchers and appear "*in black with regular quotation marks*".[6] We cross-reference quotes within the paper: quote labels appear as purple subscripts, and quote references (e.g., Q1) are clickable. All referenced quotes are listed with page numbers in an index in Appendix D.

Due to anonymity measures, readers cannot directly assess the distribution of our quotes across participants. To mitigate this, we sometimes state that certain quotes are from the same participant or from distinct participants, and Table 4 provides aggregate statistics on quotes.

Projects. Altogether, our researchers discussed 130 projects and incidents, which we refer to simply as *projects* for brevity. This category includes not only research projects involving legal risk (academic, industry, government, or independent) but also any kind of research-related legal risk incident or undertaking (e.g., receiving a subpoena for research materials, or participating in a public-interest lawsuit to mitigate legal risk in research). That said, the vast majority of projects were research projects. Occasionally, when a participant discussed a collection of research projects together as similar experiences, we treated the collection as a single project for our coding purposes. Jurisdiction was coded as that in which the researcher was based for a project.

[5] Participants were asked to optionally state their gender. We did not suggest gender categories; we recorded whatever participants told us. Participants were told gender data would only be stored and presented as aggregate statistics.

[6] The one exception to this formatting convention is the title quote, which is from a supporting professional.

Analysis. We used reflexive thematic analysis [23] to analyse interview transcripts and notes, using Nvivo for coding[24]. Initially, six full interviews were pair coded by both researchers together, to establish a shared understanding of the coding process. We divided the remaining $N - 6$ interviews between us and each interview was coded by one of us, with regular check-ins (at least weekly) to maintain consistent approaches and keep up to date on new codes. This sometimes resulted in changing, merging, or reorganising codes as we went along. Many codes required little interpretation as they captured concrete concepts mentioned by participants rather than subjective latent themes; as such, and given our work is not quantitative, we did not consider inter-coder reliability measures useful in our context [25]. After coding, we passed through the final set of codes together, performing recoding, hierarchical categorisation, and analytic reflection.

We systematically documented certain participant and project attributes, including: years of experience (see Table 1), researcher's position at the time of a project (see Table 5, Appendix E), rough timing, area of research (see Figure 1), relevant jurisdiction (see Table 6, Appendix E), type of legal risk (criminal or civil; see Figure 5), whether a legal threat was made, and whether the researcher had legal representation. For supporting professionals, we recorded jurisdiction of expertise (if applicable) and roughly how many researchers they reported they had directly supported. We do not provide complete tables of participant and project attributes due to identification risk; however, we reference this data throughout our analysis.

Where participants mentioned incidents or communications that were likely to be documented (publicly or privately), we asked whether they were willing to share the materials privately and verified the materials' contents where participants shared them (e.g., email exchanges, Twitter arguments, etc.). We used these materials for confidential verification only, not for analysis or publication; we deleted them after verification. Documents provided useful corroboration and additional detail and context (e.g., threats' precise wording). Participants generally either were willing to share the materials or did not have access to the materials (e.g., changed institutional email) and all the documents we checked corroborated the researchers' stories.

Positionality. As different people interpret data differently, and especially as we do not disclose transcripts, we explain our positionality [26]. We are both security researchers, and SP is also an attorney with experience helping researchers navigate the kinds of legal risks we study. DT is based in the UK, and SP has been based in the US and the UK; this context influenced our choice of jurisdictions to study. Our own research sometimes involves potential legal risk, but neither of us have received threats to date. We thus have some "insider status" with respect to both participant groups.

Open Science. To respect participants' confidentiality as we promised both participants and our institutional review boards, as well as to protect participants from the risks discussed in §2, neither transcripts nor our codebook can be made publicly available. For the same reasons, these materials are deleted upon project completion, so that even we do not have them. The results and analysis from them that can be shared are in this paper. Semi-structured interview questions, participant information sheets, and blank consent forms are available in [20].

2.1 Limitations and Scope

No Saturation. We originally planned to recruit 10–20 participants. In the end we had 44 participants as we kept finding aspects of experiences not covered by prior participants and new leads through previously interviewed participants. We did not reach saturation; rather, as we interviewed more, we kept learning more about the breadth and diversity of the legal risk landscape. We also found that legal risk experiences often involve aspects quite unique to a researcher's context. Every researcher we spoke with mentioned new and distinctive experiences not represented in other interviews.

As such, while our work offers unprecedented insight into a diverse cross-section of research-related legal risk experiences in the UK and the US, we stress that it is not comprehensive. We stopped interviewing to keep our work timely, and as additional interviews did not seem to bring us closer to saturation. We hope future work building on ours will further enhance understanding of the legal risk landscape.

Representativity & Numerical Results. Our analysis is qualitative, and our participants are not necessarily representative of all researchers in computer science who have legal risk experiences. We hypothesise that due to selection bias, our researcher participants may have more severe legal-risk experiences than average. That said, we are aware that there are many other researchers with at least as severe legal-risk experiences whom we did not interview. We believe it is still informative to note (roughly) how commonly certain experiences appeared *within our participants* (e.g., Figure 1). Such figures are not quantitative results from a representative sample.

Other Limitations & Scope. Our scope is limited to UK- and US-based research and to computer science (CS). Even within CS, we recognise we do not have complete representation of relevant

subfields(e.g., among fields where we are already aware of legal risk considerations, we do not have a benchmarking case or a quantum case, and for some types of research we had only one project). Snowball sampling may have contributed to getting participants doing similar work. Many of the projects discussed were years or decades ago (see Figure 1), so our data is only as good as our participants' memory; some explicitly noted they did not remember certain details. Moreover, participants may have refrained from discussing certain details even with us, due to the legal risk of doing so (see 'Discussion under §2). Lastly, we focus on *legal* risks; while we heard a little about other important risks that researchers (and their attorneys) face, including harassment and ethical risks (see §3.9.5), these are largely outside our scope (but see §3.9.5).

2.2 Ethics and Confidentiality

This study had significant potential risks for participants as they discussed legally risky activities. We developed mitigations to minimise this risk and obtained IRB and institutional ethics approval in the US and the UK (respectively).[7] Some of these mitigations were unusual and went beyond standard practices for interview projects. Our mitigations and institutional processes resulted in lengthy participant information sheets and consent forms that discussed legal risk and data practices in detail (provided in [20]).

All participants affirmed they had read their participant information sheet carefully, had an opportunity to ask questions about it, and consented to participation verbally or in writing before their interview. We offered verbal consent to avoid documentation of participation, and ultimately collected mostly verbal consent.[8] We have since deleted all consent forms.

We sent the participant information sheet in advance and encouraged participants to take as much time as they needed to read it carefully at the start of our call if they had not already done so, as well as to ask any questions they had about the information in the sheet before starting the interview. This occasionally took up non-trivial amounts of scheduled interview time and we took care not to rush participants. We took care not to seem to expect participants to have read the sheet before the call, so that they would feel comfortable taking their time reading it on the call if needed.

Recording.Audio recording was optional, and almost a third of researchers opted out. In these cases, we checked whether participants were comfortable with note-taking (they all were), and used our notes for analysis in place of a transcript. Where we did audio record, we recorded locally (not using cloud recording features such as Zoom's) and deleted recordings after transcription.

Data Security.All interviews were end-to-end encrypted (E2EE) on Signal or Zoom, or in person, except where participants (rarely) requested another platform. We stored all participant information E2EE in Signal or Keybase, or encrypted locally, with the exception of email.[9]

Outreach & Email.We determined that limiting our outreach to E2EE channels would have impractically restricted our ability to conduct the research. We preferred E2EE channels where possible, and reached out to some potential participants by email, where we did not have a more secure contact method. Due to concern about email exchanges with us heightening risk for potential participants, we emphasised in our outreach email that they should feel free not to respond, that we would circle back once, and then stop contacting them about this if we did not hear back. Where participants expressed a preference for email over E2EE communication, we honoured it. For some whom we perceived to be more at risk, we reached out only in person, leaving no digital traces.

Legal Risks. Our participant information sheets included information about potential legal risks of participating in our study, suggested participants consider consulting a lawyer before deciding whether to participate (some did), and invited participants to bring a lawyer to the interview (nobody did). We provided a one-page FAQ on Lawyer-Client Privilege and another on NDAs as part of the participant information sheet. That material was written by an author who is an attorney. In our participant information sheet and in the introduction to the interviews, we reminded participants that we are not their lawyers and cannot provide legal advice.

Legal Privilege. When requesting documentation related to legal-risk incidents, we were cognizant that some of those materials may be protected by attorney-client privilege, and that their voluntary disclosure to us could, in certain circumstances, be interpreted as waiver of that privilege. We provided participants with information sheets about attorney-client privilege and waiver in both the UK and US contexts, and when discussing attorney-client communications in interviews, we reminded participants explicitly to be careful about disclosing any privileged material to us.

[7] NYU IRB number: IRB-FY2024-9004.

[8] When discussing written vs. verbal consent, participants expressed their view that subpoenas and other tactics aiming to deidentify research participants have become more common in recent years. We encouraged verbal consent for subsequent participants in light of this risk.

[9] This required extensive justification to our institutional review boards, whose default policy prefers institutional cloud storage and institutionally approved software (which does not include Signal and Keybase).

Discussion. Notwithstanding all these measures, we are aware of cases where people were not willing to talk to us because the risk was too great. We highlight that there are serious risks that we are unable to document as the risk to participants of speaking with us is too great; this is a limitation of our study, and underscores the gravity of the legal risks researchers face. Relatedly, some participants requested that portions of their interview be off the record and redacted from our audio recording, transcription, and/or notes. Those details are more worrying than that which we have been able to use in this analysis. Participants also vetoed the use of some of the most poignant quotes about difficult situations, due to concerns about deidentification and legal risk.

3 Analysis

Our analysis proceeds in subsections, largely corresponding to high-level code categories. Each subsection is organised by sub-categories and themes.

§3.1. Big-picture themes, trends, and patterns.

§3.2. Impacts on research and researchers.

§3.3. Strategies & precautions to mitigate risk.

§3.4. Researchers' **experiences with lawyers**.

§3.5. Institutional support for researchers.

§3.6. Jurisdictional variation between UK/US.

§3.7. Areas of research that raise legal risk.

§3.8. Ethics review of legally risky research.

§3.9. Other notable themes.

3.1 The Big Picture

This section highlights key overarching themes that recurred throughout our analysis (§3.1.1), areas of research and trends over time (§3.1.2), researchers' motivations and awareness of risks (§3.1.3), and broad legal aspects including areas of law, legal threats, and legal representation (§3.1.4).

Our discussion of *trends* is largely based on participants' experiences and impressions of trends. In some cases, we also make note of patterns within our participant set as *trends*; in such cases, we note it. The trends we describe are not quantitative results, as our methods are not quantitative.

3.1.1 Legal Risks Are All Around Us

Three prominent themes recurred throughout our analysis, which we discuss next: (1) the ubiquity of legal risk in certain areas of research; (2) the importance of pursuing research notwithstanding legal risks; and (3) that nobody has a complete picture of the research-based legal risk landscape.

Ubiquity of Risks. The ubiquity of legal risk was a prominent recurring theme. Twenty-three participants[10]described legal risk as ubiquitous, unavoidable, or very common in their (clients') research fields, and were troubled by this state of affairs. See also Q1. One researcher said of their collaborators, "*I think we have all been subject to threats from vendors*". Another told us, "*I don't want to live in a legal and ethical grey zone.*" The next two quotes are from yet another two researchers.

> "*As a cybersecurity researcher, I have always been aware that we tap dance on the edge of the volcano, in the sense that it is to a certain extent our job to do things that are on the edge of what is legal and ethical, because it's our job to find attacks, and to study the way systems function.*"

> "*We had [a lawyer] on retainer . . . because, I mean, if you do forensics long enough, it doesn't matter how good you are, you're going to get sued, because it's inevitable when you do forensics.*"

But not all participants worked in fields where legal risk is common. Three researchers described being confronted by one-off legal risks and/or threats when these are atypical in their line of work. For one, their first legal risk encounter followed over 30 years in cybersecurity research. These participants expressed awareness that others in cybersecurity regularly encounter legal risksin their work, and emphasised how stressful legal risk can be for those unaccustomed to it.

Important to Do It Anyway. Another conspicuous recurring theme was the necessity of pursuing research despite potentiallegal hazards, emphasising the broader societal risksincurredif such investigations are abandoned, and often conceptualising research as an essentialpublic service (see §3.1.3).

> "*[T]his is how we make systems better.*"

Nobody Has The Big Picture. Participants consistently highlighted a systemic lack of visibility regarding the scope, nature, and impacts of legal risks. Participants emphasised incentives against discussing legal-risk experiences openly because doing so might heighten risk exposure. Even supporting professionals who helped hundreds directly and thousands indirectly believed they lacked important aspects of the big picture.

[10]Including both 18 researchers and 5 supporting professionals.

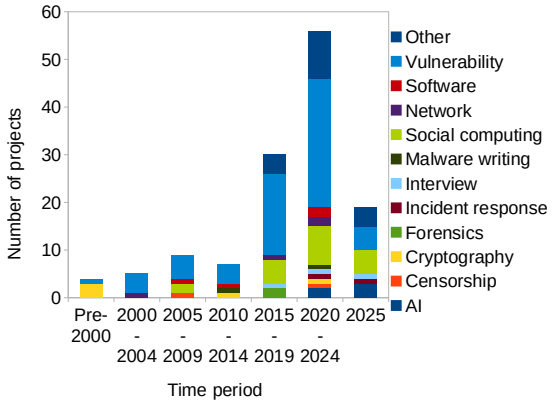

(a) All projects

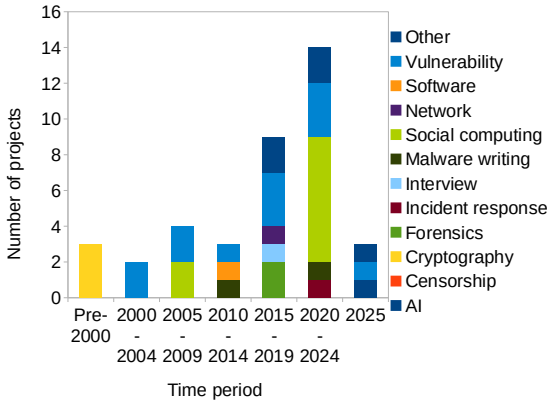

(b) Projects involving a legal threat (defined in §3.1.4)

Figure 1: Number of projects of each research type per 5-year period. Larger versions of both plots are in Appendix H. 2025 is only a part year: interviews took place late 2024 to early 2026. For readability, 'Other' bundles together usability, reidentification, physical security, and other infrequent categories.

> "*It's one of these things that you have discussions [about] . . . behind closed doors, but no one talks about publicly.*"

> "*I just feel like the whole industry is in a gray area, a gray fog, and we need to talk about this more as a community.*"

Our work makes a significant advance towards addressing this gap. See also Q8.

3.1.2 Research Areas & Trends Over Time

Our researchers discussed projects and incidents spanning over a dozen types of research. Figure 1 lists these research types and plots them over time. Most of the projects we heard about are from the last decade, with vulnerability research being a key area since 2000 and social computing surging in prominence since 2020. AI research and testing is another area that has appeared since 2020, coinciding with the explosion of interest in generative AI since 2022. Most pre-2000 projects relate to cryptography. Figure 1 also shows the wide range of other related types of research with distinct risk profiles that are observed sporadically.

Next, we overview trends over time related to: project density; research community knowledge; industry maturity; and the US political climate.

Project Density Over Time. Figure 2 shows superlinearly increasing projects over time, with a notable concentration since 2020. We would expect an increasing trend, given many participants were not active in earlier years (see Table 1); and it may be that participants favoured speaking about recent projects that were fresher in memory. But the expected increase in projects simply based on participants' length of experience would be much closer to linear (see Figure 8 in Appendix H).

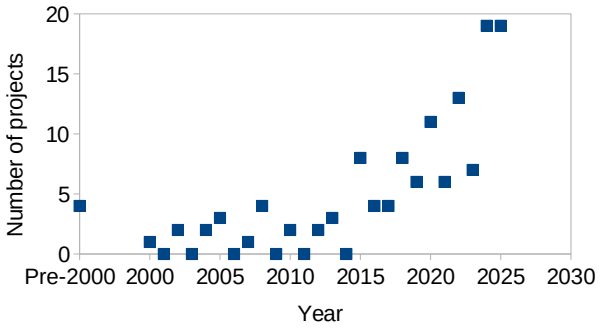

Figure 2: Number of projects described by participants per year.

We conjecture the uptick since 2020 is in part related to a surge of social computing projects facing threats (see also 'US Political Climate' below, and §3.7.2). A supporting professional's experience corroborates this hypothesis:

> « *Yes, [we're] definitely [seeing] more [threats to social computing researchers over time]. The first big change was when X sued CCDH [in 2023]. . . . The risk was mostly hypothetical until X sued, and so that changed how we advised researchers. . . . [E]ven just in the last year I probably talked to 50 or 60 researchers and journalists who worry about the legal risks they face for studying technology platforms.* »

Research Community Knowledge. Participants' emphasis and contextual knowledge often depended heavily on their main area of research. Those involved in pre-2000 *cryptography* research emphasised that the legal risks they were concerned with are all but gone, with a sea change after 2000, and that it is now generally fine to publish research on cryptography and cryptographic code.

> "*Publishing cryptography papers at all was risky once. [R]esearch[ers] in cryptography . . . are not worried about that [today].*"

Those in *social computing* discussed that the legal risk landscape has gotten worse, with some heightened risks correlated specifically with the second Trump administration, and discussed legal risk in research as a relatively recent or unusual phenomenon. And those involved in *vulnerability and network* research, who generally tended to have the most extensive exposure to legal risks in research (consistent with Figure 1), alongside awareness that thousands of researchers have faced legal risks over decades, tended to express that the legal risk landscape has gradually improved over time but significant progress remains to be made; but they were not necessarily aware of recent risks in social computing research. These categories are not disjoint, and some participants had experience with multiple research areas.

We found that community experience and anecdotes served as a primary source for legal risk awareness and identifying support networks.

> "*People learn stuff from other people in the community who have encountered similar issues.*"

However, this collective expertise often appeared to be much more concentrated within research subfields and not always connected across them. While legal experiences within a specific sub-field may often offer the most immediate relevance to a given researcher's situation, the commonalities we observe across research areas suggest that more dialogue across sub-fields could be valuable. One researcher even described connecting with colleagues outside of computer science who had experienced legal risks and offered advice they found very valuable.

One participant remarked that while exchanging experiences and anecdotes is valuable, it can also lead to the spread of misunderstandings.

> "*[P]eople chatting in the hallway . . . that's how legal misinformation spreads . . . [it's] people like me talking to their mates . . . like, 'oh, I heard there's this loophole'*"

Industry Maturity. Participants consistently told us that maturity of companies and industries — in their own development (e.g., startup vs. established company), in their security practices and expertise, and/or in having dealt with researchers before — tends to correlate with better reactions to research findings. That is, mature companies and industries were, in participants' experience, less likely to react with hostility to researchers, whether in the form of legal threats or otherwise.

> "*[I]n the early days it was very hostile. How dare you attack my product, right, how dare you impact my share price, how dare you tell me my product is not fit for purpose, and that's where a lot of the legal threats and . . . animosity between the company and the researchers came about. I think today companies have kind of matured a little where they now see . . . 'hey you found an issue that my team didn't find or maybe I didn't have budget to find, cool, alright let's get it fixed.'*"

> « *[Y]ou have to go through a cultural . . . mindset shift as an organisation to accept . . . researcher feedback. You have to go from 'we're doing everything great, look at us, we're . . . conquering the world . . . ' to this posture of basically humility in a lot of ways, 'error is human, software is hard, there's going to be bugs and it's not because we're terrible, it's because computer science is difficult.'* »Q3

Maturity may correlate with security posture. See also Q32.

> « *[M]ore mature organizations understand that actually not being assholes to security researchers, and fixing [problems], or saying that they won't fix [them] because they don't think it's is an issue, is actually in their long-term advantage* »

Mature organisations may, alternatively, simply have learnt that threatening researchers looks bad.

> « *[A]s time went on . . . the major technology companies [became] a lot savvier than [they were before], like, they wouldn't sue an undergraduate student or graduate student outright, because they sort of realized that that's a terrible look for them . . .* »

> "*I think . . . most software vendors know, even if they don't like it, that the Streisand effect kicks in pretty hard if they start publicly threatening researchers and I think that's done all of us a lot of good.*"

Participants in vulnerability research generally agreed that reactions to vulnerability disclosures have generally gotten better over recent decades, although inexperienced companies and industries still pose significant risks. That said, participants also noted that certain mature companies are nonetheless reputed for their hostility to researchers.

> « *[For] example, I['d] feel very nervous about doing vulnerability disclosure to [a particular company], because they have threatened to sue people a lot and caused all kinds of hellish stuff. That's not necessarily because they're not sophisticated, that's a very particular example . . . this company would prefer nobody knows about the vulnerabilities in [their] software or has a vested interest in . . . pretending that its stuff is secure when it's not.* »

Participants suggested that while aggressive legal tactics have declined over time, some companies continue to create difficulties for researchers in more subtle ways.

> « *[After the early 2010s,] [w]hat I started to realize was that we got threatened a ton, my clients got threatened a bunch, but it was always much more subtle, like now … there[']s a lot more soft power exercised.* »

> "*I . . . tried to do disclosure and they got very aggressive. They weren't going to sue me, but they [claimed I'd broken the law and] . . . got really in my face about it.*"

Another supporting professional highlighted the industry's growth over the same time period.

> « *[T]he changes that have happened . . . over the past twelve years have all been directionally good, but at the same time you've seen this kind of explosion of popularity [of research] that puts more stress on the system overall.* »

US Political Climate. More than ten participants discussed how the US political context, and particularly the second Trump administration, has altered the legal risk landscape for researchers. In part, the administration's shift in political priorities and personnel changes have created uncertainty around whether improvements in the practical legal risk landscape due to prosecutorial discretion and non-binding government initiatives can continue to be relied upon. Regarding people in the US government, one supporting professional remarked: « *The people I would think of as defending [researchers] are kind of gone.* »Q4 Another participant framed the issue more broadly, referring to both the US and the UK:

> « *[W]hat we have learned in the last ten years is that there's a real difference between what's a norm and what's a law. When you rely on things that are normative but not reflected in the law, and then different people with different interpretations of where the line should be gain authority or influence, you can find yourself hosed.* »

Shifting political priorities also impact researchers indirectly: e.g., through government research grants and visa policies(see also §3.9.4) , both of which were mentioned by participants.

> « *In a way that was not true a couple of years ago, if you're in the US on a visa, I would probably have more concerns about certain kinds of security research that potentially . . . violate the law now than I would've five years ago, because the legal environment for folks on visas in the United States has changed.* »

Another major theme was the emergence of legal threats that participants described as politically motivated, targeting research that powerful interests did not like — primarily in social computing.

> « *Any research relating to the spread of hate or the spread of disinformation or election interference or the like has become much riskier for researchers to engage in.* »

> "*[T]here's bizarre legal risks. A lot of it stems from people just not liking your research. We did research on topics that people don't like, and those people have acquired a lot of power these days . . .*"

This took the form of targeting by private parties, such as big tech companies, aligned with such political interests (see Q2), as well as targeting by different parts of government.

> "*[T]he federal government [has] turn[ed] the Department of Justice into a sort of aggressor going after and prosecuting people just for slighting the administration.*"

> "*It's the threat of lawsuits from people like Elon Musk [that creates] a chilling effect saying 'you don't have the right to criticise me.'*"

Referring to a situation where political pressure on a university led to the university shutting down certain research, a participant remarked, "*I thought that was incredibly cowardly. Congressmen shouldn't dictate research.*" See also Q34.

While government involvement in the legal risk landscape previously primarily concerned law enforcement (and intelligence, before 2000), participants discussed Congressional investigations targeting researchers since 2020. They characterised these as McCarthy-esque and expressed that the true goal is to intimidate and to shut down research.

> « *[P]olitics at the end of the day are driving most of it . . . the real goal is to intimidate researchers and to prevent this kind of research from happening.* »

> "*[It']s more than just a chilling effect, that is, to drown you in legal fees and make it so hard for you to do the work that you will stop.*"

Elaborating the historical comparison, another participant said:

> "*The playbook is so effective. They make very vague accusations and they don't really have evidence, the accusation is enough. The institutions fall over themselves to blacklist the people. In reading history from [McCarthy's] time, Congress never mandated people must be punished in that way. The blacklists were entirely self-imposed. The institutions did it to themselves and people need to realize that from that moment in history.*"Q5

Researchers highlighted that such politically motivated threats can be particularly unpredictable, especially given the US Congress's broad authority to conduct investigations.

> "*Things with actual legal risks, you can predict them, and you can take measures to avoid them.*"

Contrasting with prior legal risks such as in vulnerability research, one participant summed up: now, "*it's not just an angry vendor, it's the power of an autocracy.*"

3.1.3 Researchers' Awareness & Motivations

Next, we discuss researchers' awareness of risks and their motivations for doing research.

Researchers' Awareness of Risks. All participants expressed a general awareness of legal risks associated with research in security and/or social computing, including the 3 participants mentioned above for whom legal risks were atypical in their own work. Participants described such awareness as very common in their communities, for "*the whole spectrum of people from stuffy academics to scruffy hackers*".

> "*It's hard not to follow . . . closely if it looks like the government is going to criminalise the kind of research that you're doing. . . . I was in this milieu where . . . collaborators were all up on this kind of thing and so we just talked about it a lot.*"

Researchers' more detailed legal risk awareness varied much more — for example, of what activities entail risk, under what laws, and the potential consequences. One noted: "*We're not sure what bad things could happen to us.*" Another described being "*scared*" and "*super worried*" when confronted with a legal threat from a company and went on, "*What are the consequences? I don't know. . . . I don't think I ever had a clear understanding . . .*"

Such uncertainty, as well as uncertainty stemming from ambiguities in the law itself and other parties' conduct, often contributed to researchers' stress in legal risk situations: "*It's always stressful to have something like this hanging over your head and you don't know . . . how far it will go.*"Q6

Other researchers were familiar in considerable detail with relevant areas of law. Researchers mentioned they teach students about legal risks and consider that essential.

> "*I feel [it's not okay to] teach[] students to exploit buffer overflows and then send[] them out into the world without an understanding or . . . some warning of 'this is okay, this is not okay, this is a grey area in the middle that's not . . . resolved in the courts.'*"

Finally, supporting professionals shared their experience that researchers' understanding of relevant law varies based on their subfield, and researchers' misimpressions about the legality of research activities can cut in both directions.

> « *Some clients . . . I think especially scraping clients more than cybersecurity research clients, they [have] sort of a belief that 'of course I'm allowed to [scrape]' . . . and so if I tell them 'actually it's not all clear that you are', they're like, 'what do you mean, everyone does this', and . . . I agree, everyone does scrape websites, but everyone does so with the same risk that I'm telling you about right now . . .* »

> « *Sometimes [clients] think [some activity] is all illegal . . . and I'm like, no no no, this is actually pretty well protected . . .* »

Researchers' Motivations. Given such awareness of legal risks, it is natural to wonder: Why do researchers still do this work? While the law sometimes treats them like malicious hackers, participants' self-described motivations, and supporting professionals' discussions of researchers' motivations, consistently indicate quite the contrary.

Discussing a decision to collaborate with others and take on more risk, one participant recalled considering, "'*Do I want to partner with yet another*

person who's been threatened?' And we were both like, yep, because we believe in the work."

The most common perspective we heard framed research as a public service, towards improving the security of widespread technologies and/or enhancing public understanding of how societally impactful technologies work. Sixteen participants discussed concerns about problems getting fixed, impacts on users, or informing impacted stakeholders. Three recounted feeling a "*duty*", even when discoveries were made by accident, to report them for the public good, despite legal risks. Supporting professionals corroborated these perspectives. The following quotes are from distinct participants.

> « *[T]here's a lot of people that reach out kind of scared because they've found [a vulnerability] . . . and all of sudden they're like, 'someone who totally isn't me might've found this thing and now I'm worried about it, but it really needs to get fixed, and I don't know what to do.'* »

> « *[Clients ask me] 'hey I found this thing, I would like to report it without getting arrested, what do I need to be doing or be thinking about in that space?'* »

Career growth, professional reputation, and bounties were also significant motivators mentioned by 5 participants. For example, one independent researcher mentioned bug bounties as a valued source of income when they were in high school. However, in our interviews, these personal objectives consistently accompanied significant concern about the broader public good.

> "*It's not all about disclosing a vulnerability to bring attention to the researcher. I think when researchers are earlier in their career maybe they have more of a focus on that . . . but it's really important to understand that . . . the ultimate outcome is what are you trying to do[:] hopefully improve security . . . and raise standards, and that should then filter down to the bigger picture of the general public — people who are uneducated, people who interact with these systems but have no clue of the implications, like my [mother] who's in her 70s — that's really what we should be thinking about.*"

An industry researcher expressed concern about their findings being misused by clients, and the precautions they take. "*When I find an O day during a [client] engagement, unless it's part of the [explicit] scope [of my work], I don't tell the client, because I can't control what they'll do with it.*"

Researchers explained they personally were not motivated by career or money: "*I didn't care that much about the bug bounty. I just didn't want my information in other people's hands.*" Another said:

> "*Let's say my intent is to find vulnerabilities in a well-known product . . . I'm not doing it to weaponise the bug. I'm not doing it for commercial gain. I'm not doing it for insider trading or anything else. I'm doing it because I'm concerned that product isn't secure.*"

Note on "Bad-Faith" Researchers. Bad-faith researchers exist too, as discussed by supporting professionals. No researcher participants suggested they engaged in bad-faith activities (perhaps unsurprisingly, even though our outreach did not specifically target good-faith researchers). We did not encounter any indications of bad-faith research in academia.

The precise definitions of good-faith and bad-faith research are much debated in the community (e.g., [27, 28]). Here, we are referring to activities such as exploiting vulnerabilities to steal money, abuse personal data, blackmail people, or cause other harm. Participants mentioned examples of researchers who have gotten into legal trouble for exploiting research findings to cause harm, and consistently condemned that kind of activity as unethical and unacceptable. Participants moreover highlighted that the policy considerations around how the law treats bad-faith security researchers are complex, given the difficulty of drawing a bright line between good-faith and bad-faith research, given that the same laws that are used against bad-faith research can often just as easily be used against good-faith research, and given the broader public interest in having research findings disclosed even by bad-faith researchers. Further exploring the landscape of "bad-faith" research activity could be an interesting direction for future research.[11]

3.1.4 Legal Threats & Representation

We overview areas of law that came up, and figures within our participant set on the incidence of legal threats and legal representation.

Areas of Law. While cryptography export controls were a dominant concern prior to 2000, they became largely obsolete thereafter. The most commonly referenced area of law was contract law — most often by mentioning violations of terms of service (TOS)[12] — followed by the US Computer

[11] The definitions of good-faith and bad-faith research are fraught (see, e.g., [27, 28]). Here, we are referring to activities such as such as exploiting vulnerabilities to steal money, abuse personal data, blackmail people, etc.

[12] Another notable type of contract was non-disclosure agreements (NDAs). See §3.9.2.

Fraud and Abuse Act (CFAA), the US Digital Millennium Copyright Act (DMCA), defamation, the UK Computer Misuse Act (CMA), fraud, and data protection (in that order). Participants generally agreed that risks under the CFAA and DMCA have gotten better over time due to prosecutorial discretion, non-binding DOJ policies against prosecuting good-faith researchers [29, 30], DMCA exemptions, and companies gaining familiarity with security research. But they also agreed that these improvements are still not adequate for researchers to feel safe in conducting their work, and that large swaths of research remain in legal grey areas. Supporting professionals highlighted the growth of certain types of legal risk in recent years, that were not as common in the past: notably, defamation, fraud, and political legal threats.

Legal Threats. Thirty percent of projects involved legal threats described by participants. The remaining majority of projects involved potential risk without specific threats. We counted any kind of contact from governments or private parties about potential crimes or civil liability related to research as a legal threat, whether or not there was a prosecution or lawsuit. This includes cease-and-desist letters, which participants discussed as a common form that threats take. Most threats did not escalate into prosecution or litigation, in our researchers' projects and in supporting professionals' broader experience. Attorney participants shared that it is not uncommon to see threats on legally tenuous grounds that would likely fail in court.

> *« [It's] mainly threats. . . . Knock on wood, we have not been sued. . . . Sometimes it felt like we were pretty close . . . but [it's] pretty rare that we get to that level. »*

> "*I've had situations in the past where we've been sued or . . . or we get a demand letter and [our] lawyer's like, 'yeah, that might as well have been drafted in Crayola, don't even respond.'*"

Even so, these kinds of threats can still cause researchers serious stress and expense.

> "*I don't think that I'm doing anything illegal but that . . . is not necessarily the bar anymore, right, it's if you upset someone and they send a prosecutor after you, you still suffer a great cost even if you win the court case.*"

Participants discussed that there are often incentives to avoid litigation on both sides.

> *« [T]he goal is to keep things low risk, low liabilities, so we stay away from litigation when we can. It's not always possible. »*

> *« [The vendor's] lawyer . . . said 'let's write a letter . . . ' and then the client says 'I'm ready to sue' and the lawyer's like, 'that'll be a million' and the client says you never told me that. Vendor and vendor lawyer on the phone, it's happened in real time. »*

Figure 1b shows only projects with legal threats.

Legal Representation. Researcher participants discussed 30 projects for which they retained their own lawyer,[13] and 97 projects where they did not. Of the former 30, 14 involved a legal threat, 5 were industry projects, 22 were academic and 2 were independent projects. Legal representation was more common in the US than the UK (see §3.6).

3.2 Impacts on Research(ers)

> "*It affects you professionally, it affects you emotionally, it affects you in many many ways.*"

The impacts on researchers and their research are far reaching. Next, we cover in turn: researchers' feelings (§3.2.1); chilling effects (§3.2.2), and relatedly, stockpiling vulnerabilities (§3.2.3); personal and health impacts (§3.2.4); professional impacts (§3.2.5); and researchers' risk profiles (§3.2.6).

3.2.1 Feelings

The emotional impacts of legal risk are perhaps one of the most under-documented aspects of researchers' experiences. Even when researchers discuss their experiences, they often favour a factual or technical lens rather than a personal one; we found this to be participants' typical approach in our interviews as well as public documentation. When prompted to reflect on their feelings, participants revealed emotional aspects that were central to their legal risk experiences.

Fear, worry, frustration, and stress were the most commonly mentioned feelings, in that order, with 90 expressions of those four feelings across all interviews. Five researchers described their legal risk experiences as all-consuming or pervasively impacting their lives for a period:

> "*There was no escape; it was constantly high anxiety. . . . All I wanted was to have a normal day, and there was no normalcy.*"

Discussing an encounter with government agents, one researcher said,

> "*I was intimidated. My fingers were cold and my mouth was dry.*"

[13] We do not count institutional lawyers in these figures. §3.4 discusses the distinction between institutional lawyers and researchers' own lawyers.

Another researcher recalled:

> "*It was a pretty sleepless night . . . Am I going to get a knock on the door from the government in the morning, saying . . . 'don't bother to pack a bag, you're coming with us' . . . ?*"

While the vast majority of researchers described negative feelings about legal risk, a couple viewed their experiences of civil legal risk as at least a little exciting. (Nobody suggested criminal risk was exciting.)[14] One related their own such former perspective to being young and having less to lose:

> "*I was [in my 20s] so I thought it was awesome, you know . . . I had no money, so what's the possible harm? If I did it now, I'd feel very different. I own a house, I don't want to lose my house . . .*"

Another prominent theme concerned uncertainty around looming legal risk, weighing on researchers like "*a cloud over [their] head[s] constantly*" (see also Q6). Supporting professionals also mentioned this is common.

> « *It's like this terrible kind of roulette where there is legal risk associated with many common . . . research activities and that legal risk often does not materialise.* »

> « *There's never an 'okay you're good, we hereby relinquish any claims against your client' sort of thing. You're never going to get to that space. So it can be weird to . . . identify what a victory is in this space. I have to get my client ready [and tell them] it's entirely possible that at some point in this process [the company] just stop[s] talking to you, and the reason why they stop talking to you is they've decided not to make an issue out of this, and that's the most comfort I can give — which means it still hangs over your head a little bit.* »

A researcher corroborated this, discussing a situation where a company stopped responding: "*I'm sort of unsatisfied that it will always exist in limbo, but we were never sued so that's a win.*"

Finally, defiance was another notable feeling that recurred across interviews.

> "*It was more important to me to make clear where the writing on the wall was [than to save my job]. If I can say something, I will.*"

[14] §3.4.2 discusses the difference between civil/criminal risk.

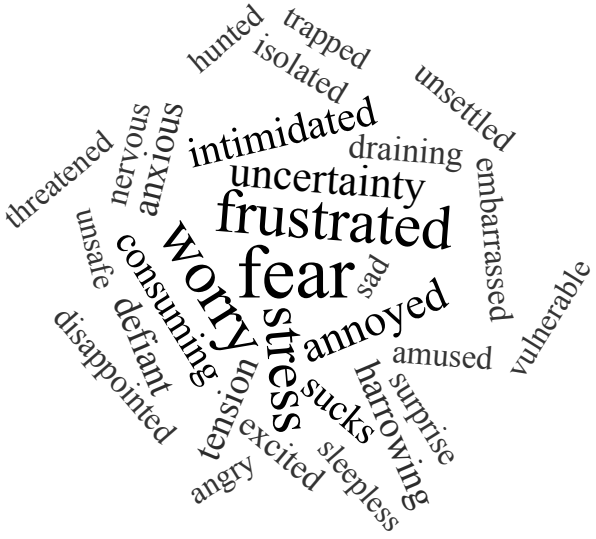

Figure 3: Feelings in legal risk experiences[15]

Feelings Word Cloud. We used Nvivo to code feelings and emotions about legal risk situations that were described by participants in our interview transcripts and notes. After completion of coding, we merged together some feeling codes that referred to very similar feelings (even if participants had not used the same words to describe them), and we aggregated a collection of codes related to fear about different things (e.g., fear of professional consequences, fear of being sued, etc.) into a parent code for fear. We then used Nvivo's word cloud generation feature to create a word cloud where each feeling code was weighted by the number of times it had been coded across all interviews. Where a given participant described the same feeling in multiple contexts, we coded it multiple times; where a given participant mentioned the same feelings multiple times in the same context, we coded it once.

3.2.2 Chilling Effects

We found that negative perceptions and experiences related to legal risks lead to broad chilling effects. Researchers discussed abandoning projects or deciding not to disclose or publish them, or avoiding certain research types or subjects due to legal risks; others shared that they have seriously considered taking such actions, and would do so if a legal risk situation escalated. The next three quotes are from distinct participants:

> "*[The risk] makes my life painful and it drives me away from [this research area].*"

> "*I just bowed and pulled the talk, simple as that.*"

> "*If risk to me, colleagues, or students is too high, I would drop [the research].*"

Supporting professionals corroborated these chilling effects, and one expressed the view that in many social computing cases, « *the real goal is to intimidate researchers and to prevent this kind of research from happening* ». The following quotes are from distinct participants:

> « *[A researcher I know] had gone through a pretty horrible time [due to a criminal investigation on their research] and told me that they didn't want to do research anymore because of the stress and risk.* »

> « *I think there are fewer researchers who are willing to do the work, and the work has slowed because it is much more costly to be a researcher doing this work . . . You have to have a tolerance for risk and harassment to do this work, and not every researcher does.* »

> "*I've kind of taken a step back. I think the threat of a company going after you legally takes a toll, so for myself I have way too much to lose now, so I'd rather let the adversaries find the bugs and exploit them, than me being threatened.*"

Other researchers, however, reported that they kept doing their research undeterred (though perhaps slowed down) by legal risks, often motivated by a conviction that the work is important for public interest, and by "*doing the right thing*". For these researchers, legal risks still impacted the way they plan, conduct, and disseminate their research.

3.2.3 Stockpiling Vulnerabilities

One of the alarming consequences of these chilling effects is that researchers are discovering security vulnerabilities, choosing due to legal risk not to disclose them to parties who could mitigate them, and keeping them around: a phenomenon we call *stockpiling*.

> « *I've seen clients choose not to disclose or choose to disclose in a much more guarded way than they would otherwise because of an apparent legal threat from the affected entity.* »

This not only creates risk in case "stockpiles" are compromised, but also means that researchers are prime targets for bribery and coercion.[16] Three researchers shared that they and/or others they know have been approached by individuals offering "*life-changing money*" to buy their stockpiles and their silence. One participant told us they didn't sell, but:

> "*I did think about it though, right, like holy shit, that's life-changing money. 'I could move my entire family out of the 'hood' money, right — so it was an ethical debate for me.*"Q7

They added that they were not sure whom the individual who had approached them represented, noting that such approaches could be related to criminal activities or different governments or even the government of their own country. "*I don't know what they use [the purchased vulnerabilities] for and I was kind of too afraid to ask.*"

Naturally, researchers' approaches to such "*ethical debate[s]*" may vary depending on individual circumstances. Two other researchers told us:

> "*I'm not at a place where I need [the money] at this point, so I am definitely more likely to disclose . . . but I can't say that that would've been the case, like if you talked to me 10 years ago, 15 years ago, I would've taken every dollar I could get, right, via blood money or not. . . . It's easy when I don't worry about my mortgage, I don't worry about bills, but if I was worrying about any of that, there's no way that I'd go to a bug bounty as opposed to a broker if I had something that could sell for what could be life-changing money.*"

> "*You spend a fortune via time and effort finding a bug for a company that hasn't really asked for it and might never pay you for it. Going to the adversary side, you find a bug in a product, you can weaponise that bug, sell that bug to other adversaries and make a large amount of money.*"

Our interviews indicate that stockpiling is happening at significant scale in the researcher community, both by researchers' accounts of stockpiling, considering stockpiling, or seeing others stockpiling, and by supporting professionals' reports of the same. Our data provides little indication of the scale at which people are attempting to buy or steal stockpiles, for whom, and how successful such efforts are.

> « *The problem with [stockpiling] is that it's kind of sight unseen; you only really realize how big of an issue this is if you're doing new studies like the one you guys are doing. It's definitely something that I[ve seen] a lot of . . .* »Q8

The same supporting professional noted that approaches may not always be overt (unlike in the case of Q7), pointing as an example to recent events

[16] See also 'National Security' under §5.

for the international vulnerability researcher community that were sponsored by certain countries with interests adverse to the US and the UK, with substantial outreach to and attendance by US/UK-aligned researchers.

Participants who discussed stockpiling generally expressed serious concern about the associated security risks. That said, one researcher emphasised that operationalising stockpiled vulnerabilities usually requires significant work.

> "*It's not ideal that there's a bunch of vulnerabilities floating around . . . but the thing that should make you feel a little bit better, hopefully, is that the process of actually making them useful is not trivial. So it's not like they have nuclear weapons, it's more like they have fissile material and . . . yeah, it sucks that they have that, you can still do dumb things with it, but it's not as bad as what you might expect [if they had] a nuke.*"

3.2.4 Personal and Health Impacts

Researchers also reported significant personal and physiological tolls stemming from the pressures of legal exposure. Describing a graduate-school encounter with civil legal risk, a participant shared:

> "*It was incredibly nerve-racking. I weighed myself before and after and I literally lost weight. I was unable to eat because I was so freaked out about what was going on. I kind of assumed that I was going to get sued into oblivion . . .*"Q9

Another participant, who had young children at the time of their legal risk encounter, recalled:

> "*I felt the responsibility that a [breadwinner] should feel about protecting [their] family, and if I went to prison, my family would be homeless.*"

On long-term impacts, another researcher, who faced criminal risk related to an undergraduate project, shared:

> "*[My experience] made me a lot more risk averse, [and I've since had] difficulty in seizing opportunities because sometimes I get really psychologically blocked on what's going to happen to me if I do this or do that, and I've worked a lot to get past that.*"Q10

More than one participant was convicted in a criminal case (these convictions were subsequently overturned), and one spent many months incarcerated. On the experience of being imprisoned, we heard:

> "*The hardest part of being jail is that there just isn't enough to do, so the perception of time is that it moves more slowly and you look for little things to break up the day . . . to try to end the monotony . . . because the more times you break up the day, the faster the time goes.*"

> "*I did get more grey hair [in jail]; it was either stress or age, I'm not sure which.*"

The personal impacts were significant when this participant was out on bail, as well.

> "*I was out on bail for a little while with [restrictions on my computer usage]. I ended up [employed in sports] for a while during this period because that was a job I could do without using a computer. . . . [T]hey . . . ma[d]e . . . sure that . . . even if I could get bail, I couldn't work because [of the restrictions on my computer usage]. So how am I going to pay for an attorney? . . . I didn't have assets, I didn't have money.*"

They told us that many years after their name was cleared, they found out they still had a felony record because the records had not been corrected, and noted they did not know whether that record had cost them any employment in the meantime.

Finally, another participant told us their collaborator is likely unable to return to their home country[17] due to their research record. Referring to a particular publication, they recalled:

> "*[My collaborator's] family member was very proud that they had published a paper, but was also in tears essentially because they knew that this means they probably weren't going to return.*"Q11

3.2.5 Professional Impacts Beyond Chilling Effects

Chilling effects were far from the only professional impacts for researchers. Six of our participants mentioned they feared for their careers, and more than one lost their job (in industry and academia) following a legal risk incident. One described "*years of being blacklisted and unemployable*" after losing their industry job. Another researcher, a graduate student when they first received legal threats from a tech company, recalled:

> "*It's just a very scary experience. There were multiple occasions where I was really quite certain I was going to be drummed out of the university and would never work again . . .*"Q12

[17] Their home country is not the US or the UK.

All but four participants kept their jobs, and we did not hear about any institutions that pushed out students. At least four additional participants discussed fearing for their careers. Participants consistently described their legal risk experiences as an intense unwanted burdenon time and resources.

> "*I was spending like twenty hours a week on calls with . . . [my institution's] general counsel, calls with my lawyer, calls with [the company]'s lawyer. It was just sucking up a tremendous amount of time. That made it hard to do work, and so that's when I think I really realised that this was an effective strategy, certainly against me.*"

Even where legal risks do not materialise, planning risky research can be burdensome. Participants shared that being careful takes longer, and could increase the risk of being scooped by less careful researchers.

> "*Everything takes forever because you're constantly running things by lawyers.*"Q13

Finally, two of researchers expressed the view that legal risk encounters can *boost* researchers' reputations.

> "*[As a graduate student, I thought that] getting sued by [company name] would be a great thing for my career. Maybe I'd be bankrupt, but I'd be somewhat famous in the industry, so that'd be cool.*"

3.2.6 Diversity of Researcher Risk Profiles

Participants highlighted that researchers with different profiles have different incentives and approaches to risk — depending on career stage, experience with legal risks, demographics, personal circumstances, and more.

> "*[Y]ou know what you can get away with, and I'm not going to lie — looking as I look, normally a bit more shaven than I am now . . . I recognise that there is the mental bias of going, 'Oh, it's fine. He's a middle-aged white guy. He's perfectly fine, he's wearing a suit, nothing to be suspicious of there.' And yeah, I can connect stuff to computers or look around the back or go into all sorts of different places and just get away with it.*"Q14

Researchers with multiple prior legal risk experiences tended to be more comfortable assessing and managing risk. Participants discussed the value of having an advisor who "*has [your] back*", and of advice from experienced colleagues. As one senior researcher put it:

> "*[Learning to navigate legally risky research is] like learning how to use a salad fork. You have to be raised by somebody who knows how to do this, and otherwise you will never do it.*"

One supporting professional emphasised the importance of communication within research teams to ensure that everyone is informed and comfortable with the risks involved, while noting that this can involve delicate power dynamics.

Three groups that stood out as particularlyvulnerable were: (1) students; (2) untenured academics; and (3) particular nationalities, depending on the research context.

Students.

> *« A lot of great computer security research comes from students. »*

Many participants expressed concern about protecting student researchers from legal risks, and those who were students during their legal risk encounters often reported "*harrowing*" experiences (e.g., Q9, Q10, Q12).

> "*I don't want my 27-year-old student to be bankrupt for the next seven years because they get sued by somebody.*"

> "*I say [to my students,] 'look, . . . you don't have to do this work, right, this is not a requirement for a PhD; we can find other projects if this is a risk for you.'*"

A few participants emphasized that students may be unfamiliar with both the legal risk landscape and norms in security research (as well as in professional adult life), which can land them in trouble they perhaps could not have anticipated.

> "*I've had students who did stupid things and almost got into real trouble and needed to be bailed out by faculty who were more senior than I was . . .*"

> "*[S]ome decisions [my students] made were very contradictory to being adult, so I feel responsible for them . . .*"

Other participants contrasted the vulnerable position of students compared to faculty.

> *« It sounds maybe obvious to say don't throw your grad students under the bus, but [it's important] to be very thoughtful about how to [protect students.] . . . Even if the PI wouldn't have necessarily done all the things the grad student did, if they gave the student autonomy . . . you've got to stand up for the folks who work for you. »*

> *« Most faculty . . . doing research, they're quite arguably an employee within the scope of employment, so there's a real argument [under agency law] that [their university] is liable, and . . . [so] the Office of General Counsel [should] handle those things . . . but there's a lot more space to argue that [institutions are not liable for] what a student is doing »*

One participant noted they protect undergraduates more than graduate students, explaining that "*[f]or graduate students . . . [risky research] is how we make systems better. Seems too risky as undergraduates.*" Finally, a participant noted students may have an expectation of support from their mentors and institutions.

> "*Lots of students participate in hackathons and you know, if any[legal risk] came up, just being told 'oh you've got to fend for yourself' was not the energy that people were looking for.*"

Untenured Researchers. Academic participants discussed the heightened risks of doing legally risky research for pre-tenure or non-tenure-track researchers. One such researcher recounted some of the most severe professional consequences that we heard about. A tenured researcher told us:

> "*[I try to] mak[e] myself the face of [the research]. If there is blowback it needs to come to the faculty. What's the point of us having tenure, if not being able to stand in front of these kinds of things?*"

Researcher Nationality. For some kinds of research, most notably around online surveillance or censorship, participants told us that researchers of particular nationalities can be at higher risk — for example, when studying online practices in their home country (see Q11). More than one faculty member told us that they have avoided getting students of certain nationalities involved in certain kinds of research due to the risk involved, and emphasized that the risks are not only legal in nature. While these risks are largely outside our initial scope — as they typically come from the legal systems of the various countries in question, and our focus is on legal risks under US and UK law — it was a notable recurring theme that we consider important to highlight briefly. These three quotes are from distinct participants.

> "*We did not get anyone to test in their own country*"Q15

> "*[It's] more than legal risk. It's a country where even if you're right . . . [y]ou want to be far enough away from it that there are no issues.*"

> "*[A particular country] has some sort of autocratic government and [researchers are] worried that their research could be perceived as going against the political norms in that country . . .*"

Another extrajudicial pressure we heard about involved concerns about certain countries stopping their citizens being students at an institution.

> "*[S]tudying China can have risks . . . some universities were concerned that if you anger China then China will no longer send students . . . and then our university will suffer . . . [T]hat made [a] university put pressure on tenure track faculty . . .*"

Risk Tolerance and Recklessness. Researchers and supporting professionals noted that those who work in legally risky research fields tend to have higher risk tolerance.

> *« There is a stereotype that has a lot of truth. People that engage in this kind of research are comfortable engaging in risk, and my job is to make sure they know what they are doing so they're taking a calculated risk . . . It's rare for somebody to listen to us talk and then say we're going to do the riskiest thing possible. Most people are sensible people who are willing to take risk. [M]y job [is] to make sure they're not taking irrational risk. »*

While participants' impression was that many researchers take carefully calculated risks, they noted that risk tolerance can also correlate with limited consideration of potential consequences, or even recklessness.

> *« Sometimes it feels like because the law is out of step with the practice, you don't necessarily need to think about what you're doing in a critical way or its consequences for people, and so I think there are ways in which that's really not great for the field. »*

> "*Some researchers clearly don't give a shit about the law and just operate under the idea that they are immune to everything. . . . [That kind of attitude] trickles down to other researchers who think, 'they did that, therefore I can do that.' So, you can kind of begin to understand why some . . . companies get a bit pissed off . . .*"

3.3 Strategies, Precautions, Advice

Next, we overview strategies and precautions that our participants mentioned can help researchers mitigate and manage legal risk, followed by a compilation of advice from our participants (§3.3.1) for researchers facing legal risk. Not all of these strategies are appropriate for all situations.

Advance Coordination. Participants reported smoother collaborations where all team members are informed about legal risks and agree about how the team will approach them, especially if such coordination takes place in advance of any legal risk incident. Individuals have differing risk profiles, risk tolerance, and incentives (see §3.2.6); participants generally expressed that it was important to take these into account. One senior researcher had a written Code of Conduct that all researchers in their lab are required to sign on to, committing to certain research processes and risk mitigation measures.

Anonymity. Participants discussed ways to keep researchers anonymous to protect them from risk. In the disclosure context, anonymity concerns were largely limited to the disclosure timeframe, and one strategy was to have an intermediary facilitate an anonymous disclosure (see §3.7.1).

In censorship/surveillance research, which can carry particular risks to researchers of certain nationalities (see §3.2.6 and §3.7.3), anonymity concerns were longer term; as such, we heard that publishing papers anonymously or pseudonymously is not an uncommon practice in this sub-field.

> "*[I]f there are risks for particular researchers involved, we try to protect them either by not having them listed as an author or by having a pseudonym be used.*"

In anonymous authorship, one of the authors is not listed on the paper at all, or in pseudonymous authorship, a pseudonym and associated email address are provided, but the author's real name is not. These measures protect researchers who are at greater risk, e.g., through nationality, not having tenure, or where their institution is afraid of repercussions (Q14). Pseudonymous authorship means that when the author applies for a job they can confidentially disclose the other papers they wrote under a pseudonym and have this confirmed by a reference letter from an author who used their real name. When publishing anonymously or pseudonymously, a participant noted that the author list must not include the word "anonymous" as this causes Google Scholar to exclude the paper from its database. Another participant emphasised these strategies are not always enough: discussing a project within their lab that they abandoned due to legal risk, they noted that it would have been easy to guess who the would-be anonymous author was, based on research context and lab membership.We also heard that similar concerns apply to acknowledgments in publications, with participants telling us that they have omitted mentioning "*people that we would have acknowledged, that expressed desire to not be acknowledged*" on legally risky projects.

Concealing research involvement can take many forms. One participant recalled a law-enforcement investigation of their non-academic research, where they refrained from mentioning a close collaborator's involvement to protect them.

> "*One of my collaborators didn't want to be targeted by the prosecutor, so [they were] happy I stopped talking about [them], until after the case was dropped.*"

Data. Participants mentioned precautions around data collection and handling, including: working with publicly available data; data minimisation; carefully thinking through who owns which research data throughout the research process; and awareness of data subjects' location/nationality and any data protection laws in their jurisdictions. Two participants noted the possibility of accidentally accessing child sexual abuse material (CSAM) during data collection targeting other types of data, emphasising that possession of CSAM is a strict-liability offence (this is true in both the UK and the US), and discussed precautions to avoid its collection and ensure researchers would not be exposed to itif accidentally collected.

In the US, for research involving collection of data having to do with health, NIH Certificates of Confidentiality (CoC) were mentioned as a useful protective measure. (But depending on context, drawing the government's attention by applying for a CoC could be a drawback; see §3.1.2.)

Contracts. Researchers reported avoiding signing any agreements or creating accounts (which can create a contractual relationship with a service provider); lawyers affirmed these can be useful precautions, especially for researchers who seek to publish their work(see also §3.7.1 on bug bounty programmes). Four researchers mentioned carefully reading TOS as a precaution, and one noted that when releasing code, data, or software, they carefully craft their own TOS to mitigate risk. One researcher mentioned carefully reading their university employment contract and institutional policies to understand how their institution would support them in case of a legal risk incident. Industry researchers noted that paid relationships with research subjects can heighten risk, again due to contractual terms.

Legal Advice. Researchers discussed obtaining legal advice as an important strategy to mitigate

risk, and four researchers reported routinely obtaining legal advice at early stages of most projects. They discussed the value of keeping lawyers in the loop throughout the research process, of consulting lawyers early on, and of having lawyers review researchers' writing. "*Don't surprise your friends, right, don't bombshell on your teammates, talk about this early.*"

Supporting professionals agreed that getting legal advice early can make a big difference.

> « *An ounce of prevention beats a pound of cure . . . in this space. If we could get to the project before it was done, there could be little tweaks that we could suggest that would still allow [researchers] to do what they wanted to do, but radically reduce their liability, or just little tactical things that we could advise on . . . . After a study's done, if there's something wrong with how the study was done, it can be a lot messier to try and solve that.* »Q16

Appendix F offers a concrete example of a research study where legal advice obtained early could enable tweaks to the research design, allowing the work to proceed with lower risk while leaving technical aspects unchanged. We heard about multiple such stories from our participants.

Institutions. Six researchers expressed that it can be helpful to socialise your research internally at your institution so that an escalating legal risk situation is not the first time that others at the institution hear about it, and so that the institution first hears about it from the researcher rather than an upset third party. Three researchers also discussed going through IRB or institutional ethics review processes as precautions that can mean institutions and peer reviewers are more likely to support the research , as further discussed in §3.8 . In network research, we heard that it can be useful to keep system administrators informed in advance of research activities.

One participant told us that while they do not give their institution advance notice of every one of their research projects, they would give institutional counsel a heads up "*where prior experience has shown that there might be some kind of issue*" or "*if we're doing something edgier that might be more likely to cause problems, then I might think about talking to the [institutional] lawyers beforehand; but for normal stuff [I don't].*"

Publication. We heard that having published some research or even having it in submission or accepted for publication can place researchers in a better position when dealing with organisations hostile to the research. See also §3.7.1.

> « *There are just frankly a lot more defenses that are in play after something gets published, and so publishing first and then preparing for the blowback [is sometimes the best strategy to protect researchers.]* »Q17

Public Relations & Public Opinion. A majority of participants stressed the importance of public opinion in shaping legally risky interactions and stakeholder responses to research, and that having a public relations (PR) strategy can help.

> "*I decided that the only way I can fight now is to drag the issue into the public as much as possible . . .*"

> "*Questions will be asked and you need to have things prepared. You need to have your public speaking voice, your talking points, your student's talk, prepare common questions and answers, make sure you have your ducks in a row.*"

In some cases, participants engaged (crisis) PR specialists to help handle legal risk situations. Some noted that academic institutions' PR operations may not serve researchers' interests, or have undermined their interests in their experience. Participants also highlighted that companies that are upset by research have sophisticated PR strategies[18] of which legal threats may be just one part, so it is important for researchers to be aware of this and prepared to respond accordingly; see Q36.

One participant emphasised the importance of « *tangible and relatable* » communication.

> « *Everybody knows what a car is and . . . why a car being hijacked is a bad thing, and so if you can talk about that . . . it's so much more evocative to a policymaker than, like, 'Let me tell you about this Citrix bug', right? . . . 'Oh, there's a guy I know who wanted to buy toys for his new niece and he tested them and it turns out they can be used by people to groom kids' — that kind of story sells itself. And then to . . . say, 'But when he tried to report it in good faith to the manufacturer, they threatened him with legal action . . . '* »Q18

We also heard that timing can be critical, and that legal threats and PR responses from other parties can move very quickly — sometimes in a deliberate effort to tell their side of the story first

[18] Indeed, we heard about at least five instances of companies impugning researchers personally and professionally in the press. We also heard about companies trying to get researchers to agree to time research announcements that might bring negative attention to the company strategically so that other news would drown it out, or after a product/version release so that they could claim the research was outdated — while dissembling to the researchers about why they were insisting on the particular timeline in question.

and to give researchers inadequate time to respond. Discussing a situation where an unfavourable high-profile news story on them dropped suddenly, a researcher described a "*crazy pressure cooker situation*" where every minute mattered in trying to get a response out to the media.

Participants opined that it can be bad optics for companies to aggressively pursue legal action against academics, and had the impression that companies have come to realize this more over time, especially in the context of vulnerability disclosure (see also §3.1.2).

> « *[It] might've been that [a particular big tech company] didn't want to see a New York Times headline saying that '[big tech company name] sues researchers.'* »

In the social computing context, in contrast, participants noted that pursuing academics can score points in the present political climate (see also §3.1.2) — and thus be worthwhile for optics even if losing money in an investigation or lawsuit that will not succeed.

Finally, a few participants adopted a strategy of silence, withholding public discussion of their research until their legal complications were resolved.

> "*Everyone I was talking to at that point, I was begging not to say anything . . .*"

One such participant later opened up publicly about their experience and considered it important for more researchers to speak out.

Research Design & Process. Researchers discussed risk-mitigating strategies in their research design and process, including: using dedicated equipment for research (not personal or professional devices); avoiding interacting with devices owned by others when possible, or getting owners' permission to do so; avoiding disrupting the active operations of any technology or service being studied; more generally, minimising collateral effects on third parties; designing research code to make clear its research purpose to observers; restricting research malware to operate only under limited conditions to avoid accidental spreading; and when releasing software, emphasising its legal uses and encouraging legal uses in practice. Participants also noted, however, that some important kinds of research questions cannot be investigated with such precautions. Conducting the minimal investigation necessary to support research findings was another risk-reducing strategy, which one researcher emphasised is important to teach students and less experienced researchers, who may otherwise take unnecessary actions that could elevate legal risk (e.g., out of curiosity while tinkering).

Seeking Permission. Three researchers discussed getting advance permission from the owners or service providers of the technology under study, as an effective precaution. In one such case, a company still sent a threatening note to researchers, but immediately backed down when presented with their own earlier written permission. But participants also stressed that getting advance permission is not always feasible, and that always requiring permission would hinder essential research:

> "*In research, asking for permission has drawbacks. [Research without permission is an essential part of] how we make systems better.*"

Stakeholders. Participants indicated that careful consideration of stakeholders who might be impacted by the research can help with assessing and anticipating risk. Disclosing research findings to impacted stakeholders in advance of publication, and giving them advance warning of publication timeline, was often discussed by participants as a risk-mitigating measure and, in the context of vulnerability research, a best practice to allow time for mitigation. Sometimes researchers shared advance copies of the research paper with companies; one researcher allowed a company to place a statement in the published paper. But another researcher discussed that they have seen that sharing advance copies of research results has led to negative outcomes including wide re-sharing without the authors' knowledge, and even attempts to "scoop" the research (see §3.7.1).

Participants also emphasised that corporate and organisational stakeholders are not monolithic, and understanding their different parts and incentives can help. Five participants shared that they have had consistently better experiences when disclosing their research to technical employees at companies and worse experiences when lawyers are involved. Some now insist on speaking with technical teams.

> "*I now ask, 'I need an expert on A,B,C,D; can you make that call happen?' Once I started doing that, things became better.*"

> "*Technical people in the companies are usually on the right side of history. They're bound by NDAs, they're bound in what they can and cannot say. One of them said [to me,] 'I'm talking to you despite a potato being shoved into my mouth.'*"

On the other hand, this strategy is not foolproof.

> "*I got yelled at by the head of the [technical team to which a given vulnerability was relevant] . . . 'why would you do this to us,*

you're a bad person' ... clearly never saw a security disclosure in their life before."

Writing. Participants frequently emphasised the value of careful writing, including: writing precisely without overclaiming; avoiding making claims about other parties' motivations[19]; emphasising precautions taken and the legality of research methods (without overclaiming); keeping a professional tone and avoiding sensationalism; and having lawyers review writing. Given that research methodologies often lie in legal grey areas, researchers suggested both clarifying details of methodology where doubt could suggest more legal risk might apply, and omitting specific details that could be used to bolster the case that the researcher should be legally liable, where consistent with research integrity.

> « *If your methods were a little risky, you're always going to have a difficult situation because at some point in the paper you're going to [need] to explain what you did, and [that] might just be a confession ...* »

Careful writing is also important outside of research write-ups, especially in the context of document requests where researchers may be asked to turn over documents they did not expect to be widely shared when they were written.

> "*[M]y emails are ... incredibly innocuous. They are boring and so there's nothing juicy to leak, nothing juicy to ask me to come in to be interviewed about, and so [the document request] just absolutely went nowhere.*"

Finally, as research can involve the handling of sensitive information (e.g., personal data, malware, exploit code, cryptographic material, etc.), researchers discussed redacting sensitive information from published research for both ethical and legal reasons. See also §3.7.1,

Other Strategies. Participants mentioned the following additional precautions and strategies: buying insurance against legal risk; keeping good records of stakeholder interactions; testing the waters with lower-stakes stakeholder interactions to gauge their approach to researchers, before discussing impactful research findings; and learning about the relevant law. Participants urged caution to avoid looking like you are blackmailing anyone: for example, avoiding language that could suggest you could use the research against them, and not asking for any reward (see §3.7.1).

[19] E.g., preferring "they did not reply for 6 months" to "they ignored us for 6 months" or "they didn't care about our email".

3.3.1 Advice for Researchers

Finally, we present a compilation of the advice that our participants would give to researchers facing legal risk. The following two pieces of advice stood out as endorsed both by researchers and by supporting professionals.

- **Reach out to lawyers early.** Supporting professionals stressed that this helps them help researchers better, and in particular, to avoid avoidable mistakes (see Q16). Researchers also stressed the importance of talking to lawyers early (see §3.4).
- **Let the public know the importance of your work for society.** Many researchers are more used to communicating in academic circles than with the public, but public opinion can really matter (see §3.3), and telling your story and expressing the civic-mindedness of researchers in a way that is relatable to the general public can make a big difference. See also Q36.

> « *[R]esearchers and their supporters need to ... make sure the public better understands what they're doing, why they're doing it, why it's important, and why the threats are threats not just to the research, but to the public at the end of the day. Researchers are engaging more and more in this way, but it's not how most researchers are trained to engage with the public ...* »

This section comprises advice that was presented by participants as recommendations or advice for others, rather than simply an explanation of strategies and precautions which they have used or heard about (which are covered earlier in this section). Some of the advice overlaps with material earlier in this section; the lists are in no particular order.

Advice From Supporting Professionals

- **Lawyers can be helpful, not just stop you doing things.** Researchers who have worked with lawyers corroborated this, mentioning examples where their lawyers suggested tweaking their research process in a way that did not make any difference to them or the research, but which the lawyers told them would mitigate legal risk. See also Q19.
- **Contract terms can really matter.** Read the terms carefully before agreeing to anything, whether for bug bounty programme, a vulnerability disclosure programme, an NDA,

or the EULA or TOS of a product or service you plan to study.

- **Careful what you tell law enforcement.** A supporting professional noted that sometimes researchers first confronted with a legal risk situation may think that an honest explanation of their good-faith research will clear everything up, but this approach can go very badly.

> *« [S]ecurity researcher culture kind of runs the gamut from people understandably paranoid about law enforcement to people who are pretty willing to believe and optimistically believ[e] that if they just make clear that this is what they did, and they were trying to help, and they were offering to fix this thing for the affected entity, [things will be fine, so] that they just sort of tell the truth about all that to the officer[, which] won't go just great for them . . . »*

- **Be aware of research community best practices.** Following research community best practices, where they exist, can sometimes help demonstrate good faith and thus reduce legal risk.

- **Do not cause collateral harm.** Be thoughtful about possible harms and design your research method to minimise them, especially regarding exposure of personal information.

> *« [R]esearchers must be very careful about causing any sort of damage or exposing personal information that is not necessary in order to demonstrate the validity of the [research]. »*

Advice for Researchers From Researchers

- **Be careful who you upset.** Whether potential legal risk becomes realized can often depend on whom your research upsets. Researchers recommended thinking through the stakeholders and who was likely to be upset, upfront, and considering any measures that might mitigate the upset your research causes.

- **Don't sign NDAs.** Seven participants' experiences were that NDAs tend to be overbroad, to constrain research, and to elevate legal risk. More than one researcher stressed that when they have refused to sign NDAs, the conversation continues anyway.

> *« Don't sign NDA or other disclosure agreements. . . . [E]ngaging in formal good faith discussion . . . and reaching an agreement in that way [is the] best outcome avoiding lawyers and legal documents. »*

- **Have a contingency plan, including how to get a lawyer.** If a legal risk situation escalates to the point that you want to get a lawyer, you may already be under a lot of pressure. Having a plan in place ahead of time can ease stress and pressure in the moment.

> *« [R]esearchers should understand the legal risks better so they can reason about them and ideally have a contingency plan in place for when things go wrong. »*

- **If it's in your means, have a lawyer review your writing.** Researchers said they found this helpful and would recommend it, but recognised that it can be costly (see §3.4.3).

- **Resist demands for information.** Whether it's a subpoena, a document request, or something else, researchers advised caution about handing over information in response to demands. Even when demands appear very official, participants told us further investigation (with a lawyer's advice) may reveal ways to either not hand over the information or to hand over less information. Participants also advised that institutional ethics committees may be able to help researchers resist such demands, in some circumstances.

- **Go above and beyond to show your good faith.** Researchers frequently discussed going above and beyond institutional or academic requirements, e.g., around data protection or ethics review, to demonstrate that you are doing your research in good faith.

- **Seek advice from other researchers you trust.** Many researchers discussed legal risk situations where they sought advice from (often more senior) colleagues who helped them through it and shared relevant experiences.

- **Take institutional lawyers' advice with a grain of salt.** Institutional lawyers ultimately have to serve the institution's interest, which may or may not align with the researcher's (see §3.4.2. Thus, participants advised caution about institutional lawyers' advice, and noted that where incentives do not align, it may be worth considering getting your own lawyer.

- **Technical and PR or legal departments can be very different.** Researchers stressed

that technical, PR, and legal departments of the same organisation can operate very differently and have different incentives, and it is helpful to keep these strands clearly separated in your head.

- **Seek the right person within a stakeholder organisation to talk to.** Participants reported very different experiences speaking with different people within the same organisation, and stressed that finding the "right" person to speak to can smooth stakeholder interactions greatly. If you realise that you're in contact with someone who is not the right fit, consider asking them politely to connect with someone with specific technical expertise.

- **Look for teaching moments when stakeholders react badly to your research.** Researchers reported experiences where a company or organisation initially reacted badly to their research, but the researcher then explained research norms and why the hostile course of action the company was considering would not best serve their interests: e.g., explaining that they would be happy to put a statement in the paper that the company was very responsive and quickly fixed the issue, which would look better for them than suing a researcher. We heard that this kind of strategy sometimes worked and sometimes did not.

- **Give stakeholders a little time to process the research.** Participants noted that stakeholders who react badly to research sometimes adjust their approach given a little time. They also pointed out that even if the researchers ultimately need to do full disclosure or otherwise escalate, having given the upset stakeholder some time looks better for the researchers.

- **Be careful with humour.** "*Be thoughtful in your humour. This is a thing that bites a lot of people — they're just joking but it doesn't matter. It doesn't matter when someone takes your email out of context.*"

- **Don't write down anything you wouldn't want to see on the front page of the New York Times.** Further yet, one researcher cautioned against putting things in writing and leaving digital traces, and recommended thinking through what would happen if any of these records leak or become public.

- **Don't expect or ask for payment unless there is an established bug bounty programme.** Seven participants discussed that this looks like extortion and elevates risk (see §3.7.1).

- **Consider buying insurance.** Six participants recommended insurance to mitigate the cost of legal risk incidents.

  > "*[T]he main risk is . . . legal cost and so if you can mitigate that to insurance, that would be something that I recommend.*"

- **Talk to the EFF.** Researchers recommended the EFF. Although they do not have the capacity to help everyone, they also have a referral network. (See Appendix A.)

3.4 Lawyering Up

Sometimes researchers expressed reservations or scepticism about involving lawyers in their research, including concerns about lawyers slowing down research or "saying no".

> « *[S]ometimes computer scientists and frankly many folks in STEM undervalue expertise that they see as kind of tied to the humanities or that is not technical, and legal work can read that way.* »

Mostly, however, researcher participants saw lawyers as an important or essential support resource, especially when legal risk situations escalate; this perspective was universal among participants with more serious legal encounters.

Lawyers themselves were keen to debunk the "naysayer" stereotype, and to describe their role as to help researchers get research done. These quotes are from distinct lawyers:

> « *Lawyers have to overcome an initial leap to conclusions you might have about lawyers, right, like we're just here to say no, we're just here to shut down anything that you want to do, if it's vaguely transgressive we retreat into our shells and then never talk to you again. . . .* ***#NotAllLawyers.*** »Q19

> « *[O]ur hope is to help [researchers] reason more concretely about their risk. We can't make a decision for them about their legal risk. They have to decide, based on their understanding of the risk and their appetite for that risk, what they're going to do* »

One of the same participants elaborated:

> *« I [see] a lot of parallels [between] computer security research [and] media law, where some degree of risk is just sort of always assumed. . . . [Y]ou'll often hear media lawyers tell newspaper clients . . . the safest thing for a newspaper to do is to never publish a newspaper — like, if you want to avoid risk entirely, don't print — but that's not the point. That's not why we're here, and so you always are accepting some degree of risk . . . »*

Four researchers mentioned they routinely sought legal advice for risky research, as a precautionary measure and for "*peace of mind . . . making sure we weren't stepping in any landmines*", even when the risks did not materialise. Others tended to defer legal engagement until a crisis is imminent.

> "*[In] every single one of my stories[, I] never had somebody come after me legally and try to threaten me over something. It's always been just caution in advance . . . and it makes me nervous.*"

It was quite common for researchers to informally consult lawyers that they knew personally or professionally, and researchers recounted seeking advice from non-lawyers that they described as legally knowledgeable.

There was also recognition that everyone having a lawyer all the time is not a satisfactory answer. A lawyer participant said: *« I [don't] like the idea that everyone should talk to a lawyer all the time; that's just not sustainable. »*Q20

Next, we discuss researchers' experiences finding and working with their own lawyers and with institutional lawyers. Then we discuss financial and resource constraints, and additional perspectives from supporting professionals.

3.4.1 Experiences With Lawyers

When you seek a lawyer. More than half of the researcher participants who discussed the timing of their seeking legal advice reported doing so very early in the projects or incidents they told us about. Other notable points at which legal advice was sought included (see Figure 4): upon receiving a complaint or threat; upon completing the analysis phase of research (and realizing based on the results that there might be legal risk); when thinking about disclosing research findings to impacted stakeholders; and when approaching publication.

Finding a Lawyer. For a vast majority of participants who retained a lawyer due to research-related legal risk, that was the first time having a lawyer.

> "*I didn't know what having my own lawyer meant.*"

> "*I was not sure what one did to [find a lawyer], wasn't totally sure what that process was . . . and who were the good lawyers.*"

On their decision to hire a lawyer, a participant recalled: "*I had no idea what else to do. Honestly . . . there was no good option, right?*" Another said:

> "*I wasn't legally savvy enough to hold my own against two federal agents that were going to interview me and try to corner me on who knows what, so it was best to hire a lawyer.*"

The most common approach was reaching out to a lawyer the researcher already knew personally or professionally. Even if that lawyer was not able to take on the work, they were often able to facilitate connections. Both researchers and lawyers mentioned that proactive outreach from lawyers to the research community can be helpful.

Lawyer Expertise. Our lawyer participants specialised in research-related legal risks, and the lawyers hired by our researcher participants tended to have this specialisation too. Researchers mostly reported being very satisfied with the extent to which they could discuss technical details in depth with their lawyers, with one remarking that this was a pleasant surprise contrary to their "*own cynical expectations*". In more complex cases, some researchers had teams of lawyers with complementary specialities (e.g., criminal defence); however, some could not afford such support (see Q23, Q24).

> "*There are more lawyers today that are much more tech savvy than . . . 20 years ago. . . . An attorney who graduates law school today, they grew up with technology, and that is a benefit even if they just have a basic understanding of how things work. That's light years ahead of somebody who grew up without any of it and is finally forced to use it so they learn the bare minimum to do whatever it is they need to do.*"

Things Lawyers Can Do For You. The most common tasks that we heard lawyers helped researchers with included: (1) reviewing writing (e.g., in the context of disclosures or publications) and (2) acting as intermediaries between researchers and institutions, law enforcement, or companies.

We also heard that the mere presence of a lawyer can improve a researcher's position.

> *« I couldn't say how many times I've had a vendor threaten a client of mine and . . . our mere appearing to be like, 'I don't think so, buddy' is enough to make them*

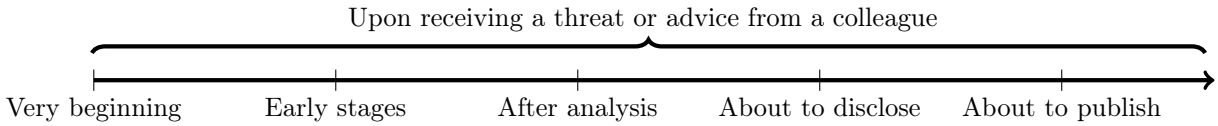

Figure 4: Common points at which participants recounted thinking about legal risks or getting a lawyer

> *go away — not always, but it is really striking . . . They on the other side might assume that this person has no access to lawyers . . . so [they] can intimidate them out of publishing or intimidate them out of doing this research, and we can appear, and then they can just back off pretty quickly.* »

> "*It was useful to have [a lawyer] around who could be kind of scary*"

Even very short-term legal engagements can be fruitful. One participant mentioned just an hour consultation on an early-stage project. A lawyer participant noted that once they get involved, a legal-threat situation can de-escalate quickly:

> « *Most of the time, they'd be yelling at us, we'd carefully point out why they were wrong, and then they went away.* »

> « *If you can convince a researcher that there are lawyers out there that are here to help, and sometimes it doesn't take a whole lot of help to see a lot of good results, I think that's really valuable.* »

But we heard that sometimes, bringing in the lawyers can also escalate tensions, or complicate the situation by prompting other parties to call in their lawyers too. One lawyer participant recounted helping a client with a disclosure while strategically maintaining the appearance that the client was managing it on their own without a lawyer.

Positive Experiences. Researchers' experiences with legal representation were generally very positive (but see Q23). Researchers generally described the legal advice they received as essential, referring to their lawyers as "*friends*" and "*teammates*". One participant expressed that they were "*very lucky*" to work with "*phenomenal*" lawyers; another shared that "*without [our lawyers,] I don't know how . . . we would have [managed].*" Researchers mentioned that involving lawyers can incur expense and delays (e.g., Q13), but still found the benefits made it worthwhile in the instances where they retained a lawyer. That said, legal advice was especially valued in higher-stakes situations, and researchers did not consider it essential to have legal representation in most research.

EFF. The EFF was the most frequently mentioned legal services provider, with 81 occurrences across all transcripts and notes. The EFF was the first support resource that seemed to come to many participants' minds. Participants discussed referrals via the EFF's Collaborating Attorneys list.

3.4.2 Institutional Lawyers

Thus far, we have focused on researchers' experiences with their own lawyers. Now, we turn to researchers' experiences with institutional lawyers: that is, lawyers representing the researcher's institution, university lawyers in academia or company lawyers in industry.

Institutional vs. Own Lawyers. The difference between institutional lawyers and researchers' own lawyers is that the former serve the interests of the institution and the latter serve the researchers' interests. Although these interests might align sometimes, where there is misalignment, institutional lawyers must put institutional interests first and need not look out for the employee's interests [6]. Researchers generally understood the distinction between institutional lawyers and their own lawyer, though often as a result of tough experience.

> "*I was like, 'the university has lawyers for us', and [colleagues more experienced with legal risk] were like, 'no, get your own lawyer'. And they were right.*"

> "*An important lesson that I learned . . . [is] that if it is more advantageous and expedient to throw you under the bus, [institutional lawyers] might go off and do [that].*"

Some researchers who had never had more serious encounters with legal risk expressed the belief that their institutional lawyers would protect them in case of an escalation of legal risk. Several UK-based researchers expressed the expectation that institutional lawyers would back up them up provided that they carefully followed all institutional research protocols such as ethics review, some believing that this was official institutional policy:

> "*I think if I'm doing the research in the university and I've gotten research ethics approval, then I would expect the university lawyers to represent me.*"Q21

Those who had had more serious legal risk incidents tended to express an understanding that institutional lawyers would not necessarily serve their interests, and discussed how and when institutional interests might align or diverge from researchers'.

> "*The thing you always have to remember is that [institutional lawyers] are there to protect the university. They are not there to protect the researcher, and that's fine, you know, someone needs to be there to protect the university . . .*"

> "*Obviously, we have personal interest and the institution has their interest and sometimes those align and sometimes they might diverge.*"

University Lawyers. Experiences with lawyers at academic institutions were *overall more negative than positive, with wide variation* between institutions. But 8 researchers recounted good experiences with institutional lawyers supporting their research: for example, sending strong responses to legal threats or complaints, backing up their researchers, which succeeded in making the other party back down; assuring researchers "*if they go after you, we will publicise it in the newspapers*" and "*don't worry, go tell them, we will back you*"; defending researchers (alongside the institution) in litigation; and supporting researchers' participation in lawsuits aiming to reduce legal risk in research.[20]

Referring to such a case, one researcher said of the decision to consult their institutional counsel, "*I am very glad I did that as I was reassured by that.*" Another shared that:

> "*I view that as sort of a defining moment in my relationship with the institution. We posed a very difficult question to them, and within 24 hours, they said yes with no reservations. That was amazing.*"

One researcher whose institution defended them in litigation told us that their university wants their research to happen, research area— even while noting that "*universities hate being sued, they avoid it like the plague.*" But another researcher whose institution defended them in litigation had an overall very negative experience with their institution and institutional lawyers. They described the institutional lawyers as effective in the litigation, because it happened to be aligned with the institution's interests, but recounted the institution seriously undermining their career in all other ways.

[20] Clarification on the last example: the institutional lawyers supported researchers' participation, but the institution did not participate in the lawsuit and the institutional lawyers therefore were not directly involved.

In contrast, one participant who has a good relationship with their university general counsel has had success collaborating closely with them on research:

> "*We have worked more closely with our general counsel [on that project] than any other project that I have been involved with so far, and it's definitely one of the edgiest things that I've been involved in. Our general counsel is in the opinion that what we are doing is fully legal and that they are prepared to defend us . . .*"

This participant recognised that their positive relationship may not be typical (Q22).

We asked several researchers whose work routinely involves legal risks whether they had considered institutional support for risky research as a factor in choosing where to take a job. One had explicitly done so, and talked to their institution's general counsel before accepting their job.

> "*[W]hen I was getting my job, I talked to the general counsel during my follow-up interviews before I accepted the position.*"

Others mentioned helpful conversations about institutional support with future colleagues who also did legally risky research, during the hiring process, and noted that the presence of others doing risky research can be a good sign. Some said they might prioritise this factor more if they move jobs again.

> "*I think I got lucky, but also maybe it wasn't an accident that I ended up being employed by two different institutions with extremely supportive general counsels towards security research. So, at [university], I mean, I was not the edgiest person [there] . . .*" Q22

Finally, supporting professionals also recounted mixed experiences with academic institutional counsel, noting that reactions vary by institution and depend on incentive alignment in each context.

> *« Sometimes the relationships [with university general counsels] are more adversarial. »*

Industry Lawyers. Researchers had *overall positive* experiences with industry lawyers and felt well protected by them. Supporting professionals corroborated that industry researchers are typically supported by in-house counsel. A few researchers mentioned as drawbacks that industry lawyers can delay the research process and publication and increase paperwork, and one industry researcher mentioned that they prefer doing types of research that

require less interaction with company lawyers, because of these reasons.

Researchers contrasted their academic and industry experiences and said they felt better protected in industry.

> "*I felt actually incredibly well protected. . . . [T]hat they were aggressively protective was refreshing, and genuinely it seemed like they wanted to help.*"

Referring to industry research at one of the biggest tech companies, one researcher expressed confidence that if any legal risk escalated, they would have a "*massive war chest behind [them] of very angry lawyers*" ready to help.

However, one researcher told us they were once sued in their personal capacity for industry research activities, a legal entanglement that took up "*two years of [their] life*", and that was the only time they hired a personal lawyer for research-related legal risk.

Another researcher told us about a project where a company commissioned research that it stood to profit from, knowing that there were legal risks involved, and accepting the risk as a cost of business.

> "*We were hired to do a job, and they were going to get sued no matter what and this was almost a business transaction, right. They knew they were going to get sued; they knew it was going to settle for a certain amount of money.*"

Finally, one researcher mentioned an unusual circumstance where industry lawyers were unable to protect their interests, when a party upset by their research findings insisted on suing the researcher in a personal capacity even though the researcher had done the work in their industry role.

Other Institutional Lawyers. Participants with government experience consistently recounted feeling well protected by government lawyers, and generally being able not to worry about legal issues with the confidence that the government lawyers would take care of any problems that came up.

> "*We were just sort of told you have carte blanche, and we just kind of trusted the lawyers to not screw this up.*"

We heard that counsel at nonprofits vary in their approaches to research-related legal risk. When a nonprofit's reputation and funding is at stake, they may pre-review all research-related communications; in one such instance, a researcher told us the lawyers were overall supportive, but constrained public messaging about their research.

3.4.3 Financial and Resource Constraints

> "*I love my lawyers. My lawyers are wonderful. They're very, very expensive.*"

The majority of researcher participants who accessed legal counsel did so at no cost from their own pockets: most commonly via *pro bono* representation (17 out of 30 projects that had legal representation), and sometimes thanks to fundraising efforts and donors. Less common sources of funding included research grants and special legal defence funds to help researchers facing legal risk.

Pro Bono. The heavy dependence on *pro bono* representation sparked frequent discussion regarding the long-term sustainability and scalability of providing such specialised legal support to a broader research community. Lawyers who offer *pro bono* services noted the practical limits to how much and what types of *pro bono* work they can take on.

> « *If it's* pro bono*, I'm spending time that I'd be otherwise able to make money and pay for cost of living.* »

> « *[I encourage researchers to reach out early but t]hat's not something that's generalizable to the entire community, because we can't help everyone.* »

Participants shared some of the factors that they consider when taking *pro bono* clients, including: the client's means; availability of other counsel; the workload and complexity of the work (including whether litigation might be involved); the public-interest stakes involved; whether the client was reckless or exacerbated the situation; and whether a client seemed easy or difficult to work with. Some participants noted that representation of difficult clients can be important in public-interest cases that will determine future law. On the other hand, being a nice human can pay off:

> « *If somebody is a total dick, I will usually figure it out fast enough that I will not then represent them.* »

Researchers also discussed the limitations of lawyers' *pro bono* capacity.

> "*[E]very lawyer I spoke to said they wanted to help me; that doesn't mean [I'll] make it onto their* pro bono *books.*"

Contingency Fee. One participant told us they hired lawyers to defend them in a civil lawsuit on a contingency-fee basis [31]. No other participants mentioned contingency fees.

Power Imbalance. Resource limitations also feed into an imbalance that researchers and their lawyers

face when confronting far better resourced governments or companies. Distinct participants told us:

> "*We're just a couple of grad students. We don't have money. We were . . . facing down a company that we knew had raised millions and millions of dollars and had lawyers and was . . . well-established.*"

> "*[T]hey were at the time one of the most powerful companies in the world, and I'm just an individual, so whether or not it violated my NDA . . . there's . . . a risk they are going to try to come after me.*"

Law Clinics. Law school clinics were a valuable source of *pro bono* representation often discussed by participants. All the work that clinics do is *pro bono* and involves law students. This means clinics have specific constraints, including limited client capacity and slower pace due to the student involvement and educational objectives, and being a better fit for matters that fit in semester timeframes.

Legal Defence Funds. Since 2023, two legal defence funds were created to assist security and social computing researchers navigating legal risks: the Security Research Legal Defense Fund (SRLDF) [32] and the Internet and Platform Researcher First Amendment Support Fund (IPRFASF) [33]. Researchers facing legal risk can apply for funding. Both are US-based, but the SRLDF is explicitly global and its grants to date have been awarded outside the US (to researchers in Malta and Poland). While the SRLDF lays out its policies and past grants on a dedicated website [32], we were not able to find much information online about the IPRFASF at the time of writing, beyond a press release [33].

The funds are only 2–3 years old, and have only awarded a few grants so far. Still, several participants mentioned these funds, with one noting, "*It's not a lot of money, it would pay for a small amount of consultation, but it['s] better than nothing.*"

Research Grants. Two researchers mentioned budgeting for research-related legal expenses in grants. Another did not believe their institution would allow research grant money to be used for legal expenses. All the researchers and supporting professionals who discussed the possibility intimated that it would be helpful to be able to use research funding for research-related legal expenses.

> "*I heard recently from a colleague at a conference that you can budget having your own lawyer into grants. Oh wow, I'd never thought of this, this is amazing. For future grants I apply for I'd like to put something in there.*"

But, on the other hand, "*[e]very dollar that is spent on something else is a dollar that's not spent on research.*"

Out of Pocket. Some participants expressed willingness to pay substantial sums out of pocket for legal representation as a last resort.

> "*Quite likely [if a legal risk situation escalated,] I would just pay a lawyer out of my own pocket, to have that peace of mind, and then maybe try to get reimbursed from the university, and if they say no, then so be it.*"

> "*[$15,000] is well worth it . . . I would pay that just to not spend a weekend in jail, so it's fine.*"

Of course, not all researchers can afford this. Both researchers quoted just above were faculty (in the UK and the US). Others who received *pro bono* support told us they could not otherwise have afforded a lawyer. And an industry participant who faced a serious criminal prosecution recalled:

> "*I never really considered hiring my own criminal defence lawyer because . . . I didn't have money. . . . I just took the public defender because my estimate was that it was going to cost in excess of $50,000 to defend myself and I didn't have that; my parents didn't have that.*"Q23

They added they were disappointed with the public defender, who "*didn't understand the technical side of the case*"Q24, noting "*while he understood the law, if you don't understand the technology, it['s hard] to apply the law to the facts.*"

3.4.4 More From Supporting Professionals

Supporting professionals shared that researchers with prior legal risk experiences and those who have worked with a lawyer before come with some knowledge and experience that can smooth subsequent interactions with lawyers and facilitate more informed risk-taking in subsequent research.

At the same time, they noted that researchers can sometimes be overconfident about their knowledge of relevant law.

> « *Sometimes I think folks can see the law as sort of a deterministic system — like, 'I just need to read the statute and I know exactly what it means and then I follow the statute' — and . . . there's a risk there where folks believe that their level of understanding is more significant [than it is,] because it feels like law is something that should work in very clear ways, [but] usually it doesn't.* »

> *« [Sometimes] it can be difficult to convey . . . that actually, the stuff you learn in law school does not get conveyed from reading a lot of debates on Hacker News. »*

3.5 Institutional Support

Institutional lawyers, discussed above (§3.4), are one aspect of institutional support. Complementarily, here we consider institutions broadly.

3.5.1 Institutional Support in Industry

Almost all researchers with industry experience felt well supported by their industry employers. It is common in industry for projects to be undertaken by researchers either because the company told them to, or with the company's explicit preapproval. Contrasting this with the academic context, distinct researchers noted:

> "*Everything has to be checked by legal—that's probably a big difference between universities and companies.*"

> "*There [i]s an incentive [for] the company [lawyers] to protect me . . . because having your employees sued for things you told them to do [is bad for the company].*" Q25

However, we had two participants who had bad experiences with industry employers, and their stories involved some of the worst consequences that we heard about. These two participants accidentally discovered and revealed problems in their own company's products, services, or practices (in separate, unrelated incidents). Their employers did not task them with that kind of research — so Q25, for example, did not apply. Both participants made their findings known internally, and when their employers failed to address the problems despite their prolonged efforts, they made their discoveries known to the public, to enable awareness and/or mitigation of the problems. Both expressed the strong sentiment that it would have been unethical to do otherwise, and each stood firmly by their decision despite the consequences they suffered.

Both of these participants left their jobs soon after their respective incidents. Between them, they faced criminal prosecution, civil lawsuits, and reputational attacks in the media and in their professional networks for many years afterward.

In the context of vulnerability research, we heard that companies with sophisticated security teams may have internal systems to handle these kinds of incidents, which would enable an accidental vulnerability discoverer to internally search for and contact the right team to handle the report, or submit a bug via a bug tracking system. Such measures are helpful, but not failsafe; one of the two participants above told us they submitted their report via an internal tracking system, and that report was grossly mishandled. In the context of social computing research, our findings suggest that even mature companies may lack established internal processes to process research findings (see also §3.7.2).

3.5.2 Institutional Support in Academia

Participants' accounts of academic institutional support varied widely, suggesting that different institutions have different approaches, and also that institutional approaches evolve over time.

We heard that in general, academic institutions are less supportive than industry employers, although we also heard about supportive academic institutions. Participants described the incentives at play:

> *« [Academic institutions] may see faculty as being employees, but they're not behind faculty in the same way [as in many other employer-employee relationships]; their interest first and foremost is the long-term health and security of the institution, which tends to eclipse the interest of any given faculty member . . . »*

> "*I don't think it's in the interest of the university, in situations like this, to actually support researchers and say 'okay, yes, we think [you have a valid legal defense].' The easiest thing for them . . . is [to] go, 'you're going to have to pay the fine.'*"

One of the most negative researcher perspectives on institutional support follows:

> "*Terrible. Terrible in every way. Maybe there are institutions that actually back up their researchers; I don't know, I've never been at one.*"

Similarly, others told us their institutions sought to avoid bad press or please powerful industry or government stakeholders by distancing themselves from and censoring research (see Q34).

But as noted in §3.4, other researchers have felt well supported by their institutions, including their institutional counsel. One suggested that while their institution is fundamentally supportive, institutions have to manage some conflicting incentives.

> "*They love the research. They would love not to have the financial and time drain of having to deal with these things*"

Another researcher's experience indicated that institutions may be uncertain how to deal with legal risk situations and whether to back up their

researchers, especially when a legal risk situation arises or comes to their attention suddenly. In one such case, a university arranged an investigation, which resulted in researchers being interrogated at length by investigating attorneys.

> "*[T]here was a lot of concern from the university as to whether or not we were the bad guys.*"

In that instance, the researchers put a lot of effort into telling their side of the story, under significant time pressure, to counter the narrative from a company casting them as "*the bad guys*". In another case where an institution was uncertain how to respond to a legal risk situation, it assembled a committee to study the situation, which eventually resulted in strong support.

> "*A very supportive [university] committee said research could and should proceed and we had an obligation to develop the technology and make it available for people [and] to the public.*"

The researcher in question added, "*I didn't feel pressured or have to change the research in any way.*" They also expressed the view that their institution has gotten less strongly protective of researchers in the years since that project.

One researcher recounted terrible institutional support at a prior institution and much better support at their current institution.

> "*[My new institution has] been very supportive. When I was about to publish my first paper in the new job, I asked if they wanted to send it through legal review and they looked at me like I was crazy, but I'd gotten so used to doing that at my last institution. They said, 'no, you're fine.'*"

Given such wide variation in participants' experiences, it would be valuable to gain a bigger-picture sense of patterns between institutions and over time. However, we found that even supporting professionals who had helped hundreds of researchers did not have enough examples to discuss such patterns, even anecdotally, underscoring the difficulty of rigorously documenting the broader landscape of challenges researchers face.

> « *My sense is [that institutions are often supportive, but] I guess I don't have a huge sample size here . . . [and] it can be that the interests are not perfectly aligned.* »

One participant described a situation where a company they did research on was an important funding source for their institution, and the institution chose to prioritize the funding over protecting the researcher.

> "*[My institution] had a special relationship with the company that was tied up in this, so . . . they would tell the company that basically it was me and it wasn't them*"

A notable way in which academic researchers receive legal threats is through complaints sent to institutional upper management such as a provost or dean (and not to researchers directly).[21] We heard about many such instances. While institutions' reactions sometimes involved doubts around "*whether or not [the researchers] were the bad guys*", institutions typically at least checked with researchers before deciding on a course of action. Some institutions defended their researchers in such situations; others shared that although their institution did not shut down their research, they did not "*get a lot of proactive protection*". Participants shared that making university stakeholders aware of the research and possible complaints in advance, and having your own lawyer, can be helpful for such situations (see also §3.3).

> « *[W]hen they perceive you as being vulnerable, they'll try and go after the people who they think have power over you . . . sometimes your lawyer's job is to prepare those people or explain to [them] that what you're doing is . . . very normal, very important, and also not something that they should be concerned about.* » Q26

We also heard of several instances where upset parties reached out to conferences at which research was being presented, to try to get the conference to censor it—including academic conferences, hacker conferences, and TED. In the instances we heard about, the conferences backed the researchers.

Finally, participants noted that researchers at public institutions can be more vulnerable than those at private institutions.

> "*[A] private institution is more shielded from the tools that can be brought to bear against a public institution.*"

Discussion. Given the unreliability of institutional support, and the possibility of conflicting incentives between institutions and researchers, supporting professionals highlighted the value of legal support resources *not* linked to institutions.

> « *I credit [university name] a lot for being pretty open [to the idea] . . . that for their students to be protected, there had to be*

[21] We did not hear about any industry researchers receiving threats in this way. This seems possibly related to a broad awareness of generally stronger institutional support in industry (see §3.5.1). Perhaps relatedly, academic researchers may more often be perceived as "vulnerable" to this type of coercion (see Q26). See also §3.5.3.

an avenue for them to get help that wasn't [through the university, and backing external support resources for their students]. It's almost the same justification as . . . an ombudsman: I'm intentionally surrendering some authority to control what this person does, because doing so builds a lot more trust in my community.»

Noting that external legal support resources are more often available for commercial ventures spun out of academic institutions, the same supporting professional lamented:

> «*[U]niversities are so willing to support for-profit entrepreneurship . . . and you know, the truth is that start-ups have atrocious batting averages, like, I've represented . . . well over a thousand of those sorts of start-ups . . . [and] I think maybe five of them are still around. . . . [Y]et the academic work that I got looped into, almost all of it had big impact. . . . The batting averages for the research . . . are just considerably better than the start-ups . . . [yet] there's big institutional infrastructure to support start-ups at universities, but [little to] nothing for independent research.*» Q27

3.5.3 Note on Independent Researchers

Independent researchers do not have institutional support as their research is not affiliated with an institution at all. Despite our participants' mixed reports of their institutions' supportiveness, we also heard that the very fact of institutional or academic affiliation can create an aura of legitimacy around research, and the mere theoretical possibility of institutional support can sometimes serve as a deterrent to legal threats — as such, independent researchers can be especially vulnerable.

> «*[Research i]s not all coming from universities, it's truly coming from everywhere, you know, the skills are pretty well distributed amongst a pretty wide variety of people in the public and sometimes like critical insights just come from people who have no affiliations whatsoever with . . . universit[ies] or anything like that.*»

3.6 Jurisdictional Variation

Overall, we did not hear about many stark differences between experiences of researchers between the UK and the US. Perhaps relatedly, we found UK researchers not uncommonly consider and navigate risks under US law. Next, we discuss jurisdictional differences in several dimensions.

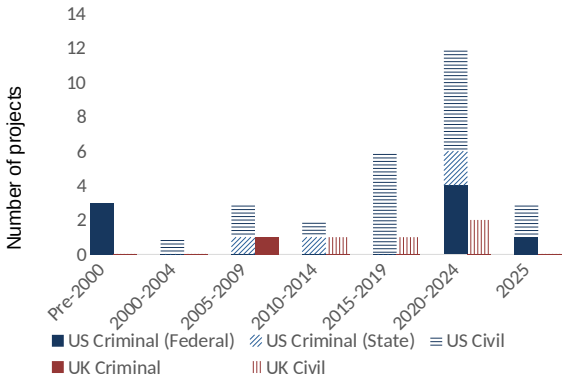

Figure 5: Criminal/civil legal threats by jurisdiction over time. Civil legal threats have horizontal or vertical stripes.

3.6.1 Legal Threats & Representation

Only 5 of 36 UK-based projects (15%) involved an actual legal threat (not just potential risk) where as 32 of 84 US-based projects (40%) involved a legal threat. Half of all legal threats were US-based civil threats (19 out of 38). The US reputation for litigiousness [34] may be relevant context for some of these figures.

Figure 5 shows the distribution of civil and criminal legal threats over time by jurisdiction. We distinguish US federal, US state, and UK for criminal threats and just US/UK for civil threats.

Perhaps relatedly, we found specialised legal services for researchers facing risk are more common in the US (see Appendix A). Having legal representation was also more common among US-based researchers: only 11% (5 of 36) of UK-based projects had a lawyer, compared to 30% (25 of 84) of US-based projects.

3.6.2 Types of Research

Overall, the kinds of research that gave rise to legal risk were similar across the UK and the US. We heard about one exception:

> "*Jamming is illegal in UK . . . We['d] never get approval for [this specific project] in the UK, but I did it in the USA because in USA there were no such rules.*"

3.6.3 US vs. UK Computer Crime Law

A key difference between the US CFAA and the UK CMA, noted by participants, is that the CFAA creates civil as well as criminal liability[22]: that is, private parties such as companies can pursue legal

[22] See [6, §2.1] for an introduction to civil vs. criminal liability.

action for CFAA violations[23], not just the government.

> « *The UK law does not include civil causes of action, and it's really hard to find examples of people in the UK who have been threatened with criminal legal action for good-faith security research in the last 10 or more years.* »

Figure 5 is consistent with the quote above. More generally, our data support the impression, discussed by several participants, that criminal action related to good-faith vulnerability research has declined noticeably over time — and that legal threats from private parties have declined too, though perhaps to a lesser extent. Our participants suggested that these trends are due to government discretion and growing security maturity in the private sector, not improvements in law. Participants found that companies inexperienced with security still often react badly to disclosures (Q30).

Our data also support participants' accounts that threats to social computing have been escalating since 2020. Figure 5 shows an uptick in US federal threats since 2020; all but one of these involved social computing. That last one, a vulnerability project, is an outlier in terms of federal targeting of vulnerability research in our data, raising questions of whether political changes were a factor (see §3.1.2, especially Q4).

3.6.4 Federal vs. State Within The US

Participants discussed differences between federal and state enforcement of computer crime laws (and other overbroad laws) against researchers in the US. In the context of vulnerability research, our participants indicated that federal prosecutorial discretion has generally been exercised in favor of researchers, and gotten better over time (eg through DOJ guidelines) though not in a legally binding way. Some thought that these days, any risk of criminal liability for vulnerability research is likely to come from states rather than federal.

The lone vulnerability project that appears in Figure 5 since 2020, mentioned above, is an outlier in terms of recent federal targeting of vulnerability research in our data, as well as contrasts with some participants' specific impressions of the federal approach to good-faith vulnerability research improving over time. This raises the question of whether political changes could have been a factor in this outlier case (see §3.1.2, especially Q4) and whether this could become a larger trend.

[23] To do this, they must meet certain fairly permissive conditions for having been impacted or "harmed" by the violation. [6] provides more detail.

3.6.5 Crossing Jurisdictions

Several of our researcher participants had done research in multiple jurisdictions, and noted there are important differences. One researcher remarked that moving to a new jurisdiction can be daunting due to new legal risks.

> "*I moved to the US and I had no idea how it works, who can you contact for which purposes. I didn't know and felt unsafe. What if someone sues you?*"

Researchers also discussed concern about legal risks originating outside the jurisdiction where they were based. UK researchers noted that US legal risks are a concern for them or their company. (We did not hear about concern in the opposite direction.) US and UK researchers discussed risks from other countries, especially when they or their collaborators had connections to those countries (see Q14), or when their research was on topics especially relevant to those countries.

Researchers expressed concern, for example, that if they travelled to these countries personally or professionally, they might be arrested or the legal risks might otherwise materialize. A few researchers expressed a defiant lack of concern about legal risks in other jurisdictions; in such cases, we followed up by asking if they ever travel to those jurisdictions (e.g., from the UK to the US, the US to Europe, or Scotland to England), to which the answer was invariably yes; however, these researchers tended to remain apparently unconcerned after this follow-up questioning.

Institutional Support. The belief that academic institutional lawyers would defend researchers against legal threats and litigation was significantly more common among UK-based researcher participants (e.g., Q21), and this expectation was met in practice in high-profile litigation in at least one instance we heard about. As noted above, legal threats were less common among our UK-based participants, and as noted in §3.4, those who had not had more serious legal-risk encounters tended more to think that their institutions would support them in case of legal escalation. Given this and our small sample size, it is hard to draw any conclusion.

Institutional Ethics Review. §3.8, covering institutional ethics review processes, gives additional discussion of differences between the UK and US.

3.7 Notable Areas of Research

Vulnerability and social computing were the areas discussed most, followed by network analysis and cryptography. We discuss vulnerability research (§3.7.1), social computing (§3.7.2), network analysis (§3.7.3), cryptography (§3.7.4), malware writing

(§3.7.5), interview research (§3.7.6), and incident response (§3.7.7) below. We note that vulnerability research is by far the longest subsection, due to participants' extensive discussion of issues related to vulnerability disclosure.

3.7.1 Vulnerability Research

> "*The fun of academic disclosure . . . I'm telling you, it's hundreds of hours.*"

Vulnerability research was the most common type of research. It encompasses activities that (intentionally or not) lead to the discovery of vulnerabilities. Often, participants discussed intentional discovery of vulnerabilities, with the plan to then disclose the vulnerability to the stakeholder who can fix it, and then to publish it after giving them time to fix it. One participant summed up the gist of a past disclosure:

> "*Here's exactly how terrible we believe this is. Please dear God respond.*"

Sometimes, though, researchers discover vulnerabilities by accident; for those for whom their research does not normally involve finding vulnerabilities, the disclosure process can be more onerous. See also §3.5.1.

Benefits For Security. Participants discussed the security benefits of finding, disclosing, and publishing vulnerabilities.

> "*Security research progresses by having people study how to attack systems.*"

> *« The accountability that comes with saying that you're planning on publishing a vulnerability . . . [is] a really important mechanism to make sure that organisations fix the [problem.] . . . I do believe that the information that's shared is for the good of both the security research community and . . . folks that might want to take action to protect themselves. »*

They shared that this does not always seem well understood by disclosure recipients:

> "*Vendors take bug reports in a few ways. The mature ones take it as a fact of life and kind of move on and get things done. . . . On the other end of the spectrum, they see it as a threat to their business and as a threat to their customers etc. and perhaps forget what fixing vulnerabilities is really for, i.e., to help their customers.*"

> "*We act like zero-days are only possible to be found by Neo from The Matrix, but that's just not reality. The stuff that we see on a regular basis is. . . I mean it's amazing that there aren't more problems going on as a result of this.*"

Indeed, we heard that the vulnerabilities participants discover are sometimes concerningly basic.

> "*It used three-byte RSA as an authentication mechanism.*"

> "*[T]o be honest, breaking into places is really too easy.*"

"Responsible" vs. "Coordinated" Disclosure. There has long been a controversy over the term *responsible disclosure* as a pejorative term [35], and this came up many times in interviews. Whether a disclosure process is responsible is subjective; the best way to handle disclosure depends on a variety of contextual factors and ethical considerations, and different researchers may reasonably come to different conclusions. All expressed opinions preferred the term *coordinated disclosure*, however, even those who expressed that opinion sometimes continued to use the phrase *responsible disclosure*. We hypothesise this may be because it is catchier and as it was the original term used, it has become habit. We follow participants' expressed opinion in using the term *coordinated disclosure*.

> "*The industry has tried to define all of this as 'responsible disclosure' for the last 20 years. 'Responsible' to me is a loaded word designed to frame anyone not following it as irresponsible, and to that I take massive offense because the reality is that responsible disclosure is only beneficial for the vendors themselves and is not beneficial for their customers. If we think about zero-day being used by the NSA. . . if a vendor decided to sit on that information, who is that responsible to, right? So my view is that 'responsible' is actually a spectrum and it really depends on what the situation is and what's going on. If there are APT groups exposing a vulnerability at any particular point in time, the responsible thing is to give information to people as quickly as possible so they can defend themselves and work out if they're exposed, not to go and let the vendor sit on it for 90 days and decide what they want to do with it.*"

> "*I hate the term 'responsible disclosure' because it implies that when you don't do it, that's irresponsible. I hate it with the death of a thousand suns, because I truly do believe it's the researcher's prerogative to just outright dump the vulnerability,*

> *right — the reality is somebody else published the vulnerable software. Now I'm not advocating for that, right, you absolutely shouldn't do it [under many circumstances] . . .*" Q28

Consensus Procedures for Coordinated Disclosure.

> "*I think it's one of those debates that's never going to be solved.*"

Participants often referred to a 90-day period until public disclosure as standard practice. For example, distinct participants told us:

> "*It is generally accepted practice to notify and give 90 days before you publicly disclose a flaw.*"

> "*Google's Project Zero kind of standardized that . . . you give them 90 days.*"

But participants also discussed the need for flexibility on a case-by-case basis.

> *« [E]very bug's a snowflake. »*

When fixing a vulnerability is difficult, a longer period may be called for, e.g., 6 months. When a vulnerability is being actively exploited in the wild or is likely to soon be exploited, then a much shorter period may be better, and in some circumstances, participants expressed that immediate publication may be the most ethical option (see also "Full Disclosure" below). When a bug has likely been deliberately planted by the vendor themselves, reporting to other partners or government may be more appropriate than reporting to the vendor. Participants also stressed that the reason for a delay is to enable fixing, and that reason does not apply when there is no possible fix.

> "*I think the industry, for the last 25 years, has really just gone around in circles around: do you disclose, do you not disclose, public disclosure, private disclosure, etc. [W]hen we publish things it's typically from a place of: how can we enable people to use this information to defend themselves, [not] how can we arm Jimmy . . . in his bedroom with a zero day.*"

> "*[Sometimes t]here is no fix, there's nothing to fix.*"

What participants disclosed when publishing their research varied. All participants expressed caution around publishing details that would facilitate malicious exploitation, or details that involve any personal data. At the same time, participants generally felt it was important to publish everything necessary for the public to understand the issue and mitigate harms if possible. Details that aid exploitation but not defence were often withheld, such as recovered private keys, or full exploit chain code. A fairly typical approach follows.

> "*All I did was give people enough information to know what the problem was and how to verify that the problem exists, but not enough information to actually break into an account. And there was a mitigation section saying what users could do to mitigate this until the company fixed it.*"

Disclosure During Submission Period. Many participants discussed that whether disclosure must be completed before research findings are submitted for publication, or it can be done during the submission period, is a area of active debate on which community consensus is lacking. Relatedly, participants told us that having a paper already submitted for publication can provide valuable protection against legal threats aiming to silence researchers (see also Q17), which researchers would have to forego if required to disclose before submission.

> "*I think 15 years ago it was very much accepted that you could submit work to a conference with undisclosed vulnerabilities, and you would rely on the confidentiality of the review process and then the usual vendor timeline is a 45–90 day disclosure window. . . . That is plenty of time to do an industry standard disclosure process, and there's a way in which having a paper that is already under review is protection against vendors who might make legal threats . . .*"

More recently, discussion of the disclosure process has sometimes been treated by reviewers a requirement for academic publication. In response to this trend, one participant said they had begun disclosing shortly before the rebuttal period so that they could say in their rebuttal that they had disclosed, while reducing the risk of the company taking action to prevent publication by submitting the paper before disclosure.

Full Disclosure. The practice of publishing vulnerability research findings without a delay or advance notice to the affected vendor or organisation is called *full disclosure*, and has been another source of longstanding controversy in the security research community. See [36] discussing a recent example. Participants generally expressed that full disclosure is can have negative consequences including facilitating malicious exploitation of research. At the

same time, they generally agreed that in some circumstances, full disclosure may be appropriate and necessary to address security problems. Some researchers felt that full disclosure was appropriate as a last resort after attempting coordinated disclosure; others saw full disclosure as the researcher's prerogative (see also Q28).

> « *I often get asked, 'Are you pro or against full disclosure?' I'm like, it doesn't matter, it's like death and taxes, right? If it's an option of last resort for . . . researchers and it's basically a symptom of failure of everything upstream of that, it's not an ideal outcome by any stretch of the imagination, but that's because other things upstream aren't necessarily in place.* »

> « *[The researcher] emailed, they called, they sent literally registered mail, they spoke with local law enforcement, they spoke with FBI, and . . . literally nothing happened, and so they went public. They documented all of this as part of their publication, all this outreach, and . . . went public with the vulnerability. I think it was after the typical 60- or 90-day period of waiting, and then as if like magic, within 24 hours the vulnerability was fixed.* »

One researcher recounted a situation in which their team felt full disclosure was the ethical choice and coordinated disclosure was inappropriate:

> "*We knew that the APT group currently targeting these devices . . . it would be impossible for them not to find it. They can't be that incompetent. . . . What are we supposed to do, just sit here? . . . It really was quite serious at the time. We can either sit here and watch this . . . continue for another three months and it was quite painful, or we could just tell people very calmly, 'this is bad, even if you think you patched please keep monitoring.' [I]t's a good example of where again people wanted to apply 90 days arbitrarily, but context matters and in that context [we believed] we had to let people know.*"

Bug Bounties and Vulnerability Disclosure Programmes. Bug bounty and vulnerability disclosure programmes (VDPs) have become commonplace. In a bug bounty programme, vulnerability reporters may be offered a financial reward for their report. A vulnerability disclosure programme is like a bug bounty but without financial benefits. While being paid for research findings can be beneficial for researchers, we heard that both types of programmes have problems and can be a "*massive failure mode.*"

Participants told us that bug bounty programmes can have restrictive terms about what kinds of research are permitted, which in practice may prohibit meaningful research activity or even create heightened risk for researchers who go "out of bounds." We also heard that programmes often require NDAs that prohibit public disclosure or discussion of research findings with others, prior to determining whether payment will be made.

> « *Best-case scenario is they actually have a process for [reporting security issues] and that process does not involve signing an NDA through HackerOne and then waiting 400 years. The first thing I would look for is: Does this organization have a process for doing this? If so, then can we follow that process? Are there reasons why we wouldn't want to?* »

Recipients often have an incentive to silence the reporter, rather than fix the vulnerability.

> "*VDP processes are focused on . . . how do we ensure that this stuff doesn't get out?*"

Supporting professionals advised caution in using bug bounty programmes. See also [37].

> « *On HackerOne there are some bug bounty programmes that I would never advise a client to sign on to, especially if they were going to . . . publish [their research].* »

> « *My job is to think about that, to think about leverage on some level and to say 'hey let's not have you sign something or agree to something that basically puts the company entirely in the driver's seat here, because . . . it's also another potential legal claim that they could have against you if you do disclose, right, then you violated your contract with them, and that's not a great place to be'.* »

Participants discussed going to great lengths to avoid using bug bounty programmes to disclose vulnerabilities, generally because they needed to disclose the vulnerability publicly (e.g., in a paper or DEF CON talk). This can be difficult; we heard often that vulnerability recipients try to push researchers to use the bug bounty programme.

> "*[W]e avoid all VDP processes, all bug bounty policies. We'll actively go and hunt for individual employees [to disclose to].*"

An attorney participant noted that some researchers may derive a false sense of security

from reporting their research via bug bounty programmes, as the programmes give the impression that the recipients welcome disclosures, when in reality doing so binds them to disadvantageous terms.

> « *I've known a lot of clients who have submitted [research findings via a bug bounty programme] because they . . . figured that's the safest thing to do, before talking to us — and then it turns out that the terms of that bounty are really bad for them, but now they're kind of locked in because they've . . . acted as if they're agreeing to the terms of the bounty by submitting via the [bug bounty submission] form . . . and so it just puts you in a real bind.* »

One participant lamented that submitting via bug bounty programmes meant they could not disclose the vulnerability, but that they had also never been paid for any of the bugs they had reported via these programmes.

> "*You cannot speak about those vulnerabilities [submitted to VDP or bug bounty programmes], whether they do or don't accept them. That's entirely up to them. There's no period in which you can disclose them. Effectively you are completely under the control of the vendor at that point, and that stuff we just vehemently disagree with.*"

On the other hand, we heard that some organisations welcome researchers' submissions to their programmes and honour the programmes' terms. Some participants saw bug bounty programmes and VDPs as a sign that an organisation may be friendlier to researchers, and noted that they might feel comfortable undertaking types of research that they might otherwise avoid due to legal risk, if explicitly permitted by a programme.

Even further, we heard about bounties sometimes being paid out for vulnerabilities that are clearly out of a programme's scope. One participant told us this is common enough that some bug hunters focus on out-of-scope vulnerabilities as they think there will be less competition. Researchers also reported being offered bounty payments for disclosures despite not seeking them.

Finally, some industry participants shared that just as bounty-hunting researchers need to pay close attention to the terms of bug bounty programmes, industry researchers need to be careful to stay within the contractual scope of paid engagements.

Who To Disclose To. We heard consistently that finding an appropriate contact to disclose to can be difficult and time-consuming. Participants used techniques like "*randomly guessing security@company.com*", "*hunting for people on [social media such as] LinkedIn and then . . . hop[ing] that they reply*" (a risk if fraudulent profiles are listed), trying to find an existing contact who has a contact at the company, using an intermediary such as a CERT (Computer Emergency Response Team), or using press contacts or a customer support form.

A key challenge is to get the report to someone with the expertise to understand it, who will not react with hostility or legal threats, or ignore the researcher. Distinct participants told us:

> « *The hardest part . . . is trying to figure out the right person to talk to and connect with.* »

> "*[F]iguring out who some data belongs to is a completely ad hoc, messy, unreliable process.*"

Some participants explicitly request a meeting with people within an organisation who have specific relevant technical expertise when they make the initial report, in the hopes of having one meeting where they can be sure that the recipient understands the issue — rather than having protracted discussions with a succession of people who do not have the relevant background, which participants had experienced and found frustrating.

> "*We found the email for the security team, and we were like 'hi, we're academics.' . . . They couldn't understand anything we did. We explained again and again and again, they didn't get it. Can you get the person you designed this particular system on the call? They did, we told them, and we never heard from them again. Ever again. . . . It was very unexciting.*"

However, one participant reported having unusually good university infrastructure to facilitate outreach.

> "*I just say 'okay I have this problem, I want to contact this company' and if my admin staff doesn't know the contact, someone [else] will have, because there are a lot of admin staff.*"

Unresponsiveness To Disclosure. We heard about many, many instances of recipients never responding to disclosures, either apparently ignoring them, or sometimes silently fixing vulnerabilities (possibly undermining researchers' publication of the research); the unresponsiveness of recipients came up in 16 interviews. For example, distinct researchers told us:

> "*[A company] literally ignored us for months. . . . [The] process can become very*

frustrating when trying to do the right thing. They silently patch vulnerabilities, don't tell organisations, don't tell clients, and then we're . . . left in this limbo."

"*More often than not, we don't get to talk to companies because they either don't want to talk to us or it's not even clear how to get in touch with them. It's just a black hole.*"

"*A major AI company has a basically 'take it and will never communicate with you again' kind of policy. I don't know if that's the written policy or not. I don't know, but they take it, and then you go and check again and like three days later it works, four days later it works, five days later it's fixed, right, but there's no acknowledgment that anything was fixed. It's just like firing something off into the void.*"Q29

Sometimes, we heard, recipients who ignore disclosures then threaten researchers when they try to discuss the vulnerabilities publicly. Others only respond once told about forthcoming publication.

Some recipients only take the vulnerability seriously if the researchers can prove exploitability, which can take significantly more engineering effort. Sometimes, researchers rediscover vulnerabilities years later that have still not been fixed, but are unable to determine whether they were disclosed to affected organisations by the original discoverers.

Reactions of Disclosure Recipients. When organisations did respond to disclosures, their responses varied widely.

"*[I]t could go one of two ways. The company could openly embrace that and say 'good catch, thanks very much we're going to fix it and we'll work with you' or they could go down the other approach and say 'we don't like what you've said, we're throwing the lawyers at you.'*"

A supporting professional elaborated:

« *Reacting in good faith . . . looks like them saying they will fix it, keeping the researcher informed about the timeline for that, asking follow up questions to make sure they understood the vulnerability and fixed it, all the way up to offering a bug bounty or offering to present jointly at a conference. It makes it look like they support the research community. . . . Bad faith might be stringing [researchers] along and saying we heard you, we'll get around to it or quibbling [on] what the report says, . . . like, 'we don't really think it's a big deal, our users don't care about this sort of thing.' And occasionally, [bad faith] will look like a small company taking it personally: . . . 'cease and desist, we will sue you, you've broken a law . . . and if you publish, we will hold you accountable . . . '* »Q30

Researchers shared direct experiences of all three types of reaction outlined in Q30.

"*We've seen some teams actually take the time to do major architectural changes after flaws have been found.*"

"*[The owner of the vulnerable system argued that:] it's not a vulnerability because nobody has ISP level compromise capabilities — except every nation state does . . .*"

"*[The software company] hammered us because we went to them and said 'hey, we want you to fix this: these are serious vulnerabilities that are impacting us, as well as all your other customers.' And they wouldn't do it. . . . This was a CVSS-10, like no question CVSS-10 . . . We finally were like, 'okay, we're at least going to notify some of [your other customers, whom we work with] that we know [are impacted]' . . . but they were just hardcore like 'You can't tell anyone else. Your legal agreement precludes you from even notifying other [customers of ours] about this vulnerability. So then we sunset them as quickly as we could.' . . . That was a big wake-up call for me, an enterprise grade piece of software where they not only are like we're not fixing this, but we're also not . . . we're going to legally . . . hammer you if you even notify anyone about it.*"

Yet another type of reaction participants reported experiencing involved denying the consequences of the research findings, explicitly or implicitly. The next two quotes are from distinct participants.

"*[We] reported . . . and obviously expected them to do something with it, and they wouldn't at all, and first tried to deny the vulnerability existed . . . you know the game, right.*"

"*They bury their heads in the sand, pretend the exploitation's not happening, and we just watch their clients get absolutely wrecked.*"

Intermediated Disclosure. Participants described disclosures through an intermediary such as their lawyer or a government agency (e.g., CISA in the US or regulators in the UK), especially where they were particularly concerned about a hostile reaction (e.g., due to a company's prior history). Sometimes the intermediary kept the researchers anonymous (see also §3.3). Intermediation can be useful during mass vulnerability disclosures where it could otherwise be difficult to coordinate with all the recipients. All the intermediated disclosures we heard about went comparatively smoothly.

Heuristics For Risk. Participants discussed various heuristics they used to informally gauge how risky reporting a vulnerability would be. Participants reported that larger, more mature technology companies generally understand the vulnerability disclosure process, while smaller non-tech companies can be confused and respond unexpectedly. We heard that organisations with clear vulnerability disclosure processes (e.g., in security.txt on their website) often react well, but many systems lack this. Researchers suggested small companies and startups can react well, but are more unpredictable; one researcher noted that if you report to a CEO or founder inexperienced with security, they may feel a more personal threat, then it can go badly. Another researcher observed:

> "*[P]eople react badly when they lose money.*"Q31

Relatedly, we heard that another influential factor is the vulnerability's commercial implications:

> "*[T]he problem is when you're sitting on a result that actually matters commercially, you're going to start running into adversarial actors that don't have your interest in mind. They want to see that paper gone, right.*"

Participants shared consistently smoother experiences reporting via government intermediaries such as CISA (Cybersecurity and Infrastructure Security Agency) in the US and NCSC (National Cyber Security Centre) in the UK. See also §3.3.

Participants were especially cautious about reporting to companies with a reputation for being litigious, sometimes saying they decided not to report to such companies at all; see also §3.2.3 on stockpiling. Certain industries had reputations for being particularly risky to disclose to; higher-risk industries highlighted by participants included the automotive industry, Internet of Things (IoT), finance (perhaps related to Q31) and medical devices.

> "*[Some] medical device manufacturers . . . are almost adversarial towards any security research . . .*"

In general, participants told us that larger, established companies with stronger security practices and mature security teams tend to react better.

> « *Companies that I generally think of as making relatively secure software are interested in vulnerabilities. Companies that would like to look like they're making secure software [when] they don't . . . mak[e] life hard for anyone who wants to tell them that they're doing anything wrong — those are the people I worry about.* »

But contrary to the general trend, we heard that some large security technology companies have bad reputations, and sometimes some individuals at experienced organisations may not know how to handle disclosures, even if other parts of the organisation do. Participants also suggested that older organisations with lots of legacy systems that they struggle to maintain can be riskier to disclose to.

In some participants' experience, reporting vulnerabilities to governments could also be risky, particularly smaller or regional government agencies that may not understand vulnerability disclosure. One researcher said they avoid research on military technology, due to the heightened risk.

An often mentioned factor, regardless of industry, concerned who the report reaches first: if it gets to a technical person who understands the system and the vulnerability, then we heard they are usually appreciative, whereas if it gets to lawyers first, our participants suggested there was more likely to be an adversarial reaction.

> "*The biggest problems happen when either a big company gets surprised by [research and] then the legal people appear before the technical people do, or with small inexperienced companies they can take it the wrong way and go to the police or something.*"

> "*[I]f we had just disclosed directly to the technical people, they would've been like 'oh cool, thanks' and fixed it. I mean, once we had a phone call with the technical people they were like, 'oh yeah alright cool thanks', and they fixed it within 24 hours. But it's been months and months and months of the lawyers, the company lawyers wanting various things.*"

Some participants highlighted the broader point that organisations are not monolithic: we heard about one situation where some employees are not aware that they had a bug bounty programme (or informed about what one of those is), which led to chaotic reactions when a vulnerability report arrived.

Finally, one researcher noted that in their experience, European organisations are often more open to disclosure than US organisations, perhaps as they are less concerned about being sued.

Asking For Payment Looks Like Extortion. Researchers discussed that with the rise of bug bounty programmes, it has become more common for researchers to request payment for vulnerabilities, even in situations where there is no established bug bounty programme. But requesting payment can easily look like extortion, which may in turn lead to a hostile reaction and escalation of legal risk.

> « *'Hey I've got this thing and I'll let you know how I'll totally hack you into the ground, if you pay me money' — that's not a good way to have the conversation. I think even creating the appearance . . . that that's why you're doing it is pretty risky.* »

Requesting payment when no bug bounty programme was offered has come to be known pejoratively as a "beg bounty". A supporting professional noted that those who seek beg bounties often try to turn a profit by sending out a large number of very minor findings to many companies, hoping a few of them pay out, and thus take up resources that would be better spent elsewhere.

> « *My favourite description of [beg bounties] is when you pull up at the traffic lights and the dude jumps out with the squeegee. . . . It's like, okay, it's sort of useful, I guess, but I didn't really ask for it and it's not that valuable and you're kind of annoying me right now, so stop.* »

Avoiding making money has other benefits. One participant described how not making any money from conducting the research minimised legal risk because by not obtaining anything of value, they ensured that certain fraud laws did not apply; another noted that payment tends to imply a contractual relationship, which can heighten legal risk.

Finally, participants noted that even when no payment is discussed, the standard coordinated-disclosure practice of giving recipients a set period of time such as 90 days to fix the problem can look like extortion to those unfamiliar with coordinated disclosure, which can lead to hostile reactions.

> « *[If a researcher says] 'I'm going to give you this many days,' you know, the first time you [hear] that, it can feel like a threat, so you want to be careful about how you voice that.* »

Explicitly stating that you are not seeking any payment may be helpful; we heard about one instance where a disclosure recipient initially reacted with hostility but calmed down once told the researchers were not after money.

Disclosure as a Gift. Researchers expressed the perspective that their research findings are akin to a gift or donation that helps organisations make their systems more secure, pointing out that if an organisation commissioned a security review to get the same results, it would be very expensive.

> "*It's a free pentest, right — they get their product made more secure. They don't have to pay a lot of money. In fact, they have to pay hardly any money, and yeah, it's kind of a win-win.*"

> "*If you disclose to a [mature organisation], then you're probably fine, because they have a security team that understands what we're doing is kind of like donating a large chunk of money, because this sort of analysis that we're doing would be really really expensive for them to do.*"Q32

Supporting professionals agreed that this is a common researcher perspective, but added that this perspective misses some of the broader considerations and incentives on the organisation's side: « *you're giving them a bad day even if you're doing the right thing* ». They suggested greater researcher allowance for these broader organisational considerations may help smooth disclosure interactions.

> « *The thing that I try to remind my clients is that you're presenting the vulnerability to them, and you think you're doing them this favour — and in a big way, you* are *doing them [a] favour by telling them this. But that's not the end of the risk on their side, because they're not just thinking about 'is this vulnerability out there,' they're thinking about: were any of my customers affected, am I now going to have a lawsuit from one of my customers, is one of my customers going to cancel an agreement with us because this vulnerability was out there, am I now going to be subject to an investigation by the FTC because I had bad cybersecurity practices? There's that whole other layer, and even though you're giving them a gift on one level, you are introducing a whole lot of potential trouble . . .* »

Other supporting professionals noted that it can be frustrating for both parties if a researcher's assessment of a vulnerability's seriousness does not align with an organisation's assessment, and cautioned researchers against inflating the significance of research findings. Distinct participants said:

> *« [I]f the vulnerability that is discovered is just simply not that impactful, it may be further down the queue than other vulnerabilities that the software vendor is working on and researchers I think need to keep in mind that broader context. »*

> *« [A] good way to piss off a company is to make a huge deal out of something that's not actually a big deal. »*

Cryptography. A couple of participants observed that cryptography vulnerabilities are often hard to exploit and easy to fix. As such, their disclosure, while academically interesting, may not incur so much legal risk as, for example, vulnerabilities that enable remote code execution.

AI. Reporting AI vulnerabilities is complicated by the fast-moving and immature nature of the field with, for example, academic AI venues not expecting disclosure of vulnerabilities and unaware of coordinated disclosure norms (see also [38]).

> "*I think we are in the world of the 1980s in security in ML. Nobody is talking about it. . . . This is where the wild west exists, and nobody's reporting anything to anyone.*"

Moreover, we heard repeatedly that what even counts as a vulnerability in AI systems is contentious, and recipients can struggle to understand vulnerabilities due to lack of security knowledge.

Participants consistently found Google and Microsoft more responsive than specialized AI companies (see also Q29), and noted that their being more established companies, their existing security expertise, and their extensive experience with vulnerability disclosure outside of AI may contribute to their better security posture in AI (see also §3.1.2).

> "*Outside of Google and Microsoft, nobody [in AI] cares [about security]*"

One participant shared that when they disclosed vulnerabilities to a major AI company, the company's response was that their product could not be insecure as they had outsourced the security to a general-purpose cloud services provider. They then tried to deflect the researchers to talking with the cloud provider, even though the vulnerability was in the AI product and not the cloud services it ran on. Others reported not even receiving a response.

More broadly, participants shared examples of AI products being shipped with little concern for security and systemically vulnerable. They also raised concerns that the techniques to defend against these attacks are still under development and that they have seen AI companies respond by writing disclaimers saying the products are insecure to use for their intended use case, rather than to provide security by default and fix vulnerabilities.

Participants also highlighted that AI and security are both highly specialized fields, and part of the challenge in securing AI systems lies in the scarcity of joint expertise.

> "*It's hard to find talent that understands both security and AI, and the problems are hard. There are lots of problems we don't know how to solve.*"

Yet another risky aspect of AI research is that it is often difficult to conduct without touching third-party infrastructure, meaning that it is easier for AI companies to detect research activity and to invoke TOS violations and/or computer crime law.

> "*Anytime you're doing AI testing, with very few exceptions, you're touching backend infrastructure, right, infrastructure that's doing the inference.*"

One participant preferred to study open-source rather than commercial AI models where possible, to mitigate legal risk from commercial actors; however, they noted that to understand the state of the art, it can be to investigate commercial models.

However, none of our participants reported that any AI company has gotten litigious over research activity. We heard that the AI companies do seem to automatically monitor for research and other activity they do not want on their platforms, and some of them terminate flagged accounts. These two quotes are from one participant.

> "*The AI companies right now are draconian about [users who test their systems] — one major AI company in particular. I've had two accounts cancelled by [that company] — paid accounts, disabled. [They said] I was doing unauthorized security testing against their terms of service. . . . I'm doing the same work with [two other AI companies] and they're not terminating accounts . . . [As for Grok,] I mean, no surprise they're not doing anything, because they're not monitoring anything there.*"

> "*I think the lawyers at the AI companies are pretty busy with IP issues right now . . . for my $20/month account probing this stuff, they're just like, 'terminate, terminate.' Then I'll set up another and do it again.*"

Finally, participants remarked there may be competitive reasons behind AI companies' measures to

stop research testing, as it is common for AI companies to perform adversarial testing on one another.

Other Disclosure Experiences. Researchers shared that a big technology company being the finder of a vulnerability can also make disclosure run more smoothly.

One participant shared that vulnerability disclosure can be complicated by vulnerability recipients "scooping" or producing their own papers for submission based on a disclosure, or widely circulating a disclosure and accompanying paper to industry mailing lists without researchers' knowledge. They stated that they have direct knowledge of multiple instances where this sort of thing happened.

> "*I don't want my paper circulating through hundreds of people whose names I don't even know because I might get scooped and therefore my students might lose out. I can't vocalise that without being judged.*"

3.7.2 Social Computing

Social computing was the second most prominent research type, with four times as many projects and twice as many legal threats as the next largest category. It encompasses measurement and algorithmic or platform transparency/accountability research where the aim is to measure the behaviour of people using a platform or the impact of a platform on society. For example, this includes projects scraping data from forums or social media, analysing ad campaigns, and investigating disinformation or moderation practices. While not everyone referred to this area by the same name, we adopt the term *social computing* used by many participants.

Participants in social computing had their own legal representation more often: 35% of social computing projects had legal representation (cf. 16% of vulnerability projects).

> "*We had a team of 8 lawyers who were keeping us from being sued. . . . [But j]ust think about how many lawyers [big tech company] had on this . . .*"

The earliest social computing project we heard about was in the late 2000s, and prior scholarship has documented legal issues around scraping (under the CFAA) in courts in the US since 2000 [10]. Social computing projects were comparatively rarer in our data until around 2015, and by far the highest concentration of legal threats in social computing has occurred since 2020 (see Figure 1). This trend was corroborated by participants who had been active since the 2000s.

The history of legal risk in social computing thus appears to be more concentrated in recent years. We heard from participants, including ones experienced with both social computing and vulnerability research, that community expectations and approaches to legal risk are less mature in social computing than in vulnerability research, likely due to the recency of the the surge in legal threats, combined with the unpredictability of politically motivated threats (see §3.1.2).

For example, we heard that researchers have fewer established procedures for conducting such research (cf. norms around disclosure in the vulnerability context; see §3.7.1), and that organisations implicated by research often have less experience dealing with social computing research, and fewer norms and programmes designed to encourage research. While programmes to enable researcher access to platform data exist and are valuable, for example, our research suggests that they are less prevalent and normalised than, say, bug bounty programmes, and recent advocacy has pushed for more robust frameworks [39]. Moreover, participants agreed regardless of research area that lack of experience dealing with researchers correlates with organisations reacting badly to research.

A key theme in social computing — which also came up in other areas, but was particularly prominent here — is that many projects violate Terms of Service (TOS). Moreover, we heard TOS violations can be necessary for essential research (see also Q37).

> "*[C]ybersecurity is about understanding this gap between how people understand a system to function and how it actually functions, so very often how people understand a system to function is defined in terms of service, and so if we can't violate terms of service, we can't test how a system actually functions.*"

A participant explained that research subjects' attitude towards researchers depends on context.

> "*[T]he adversary in some cases is the platform itself . . . this is an adversary that has a lot of resources and has a vested interest in refuting our findings when they make the platform look bad. There's also times when our conclusions say these guys are doing something good, and then we're not an adversary, we're not in that position.*"

The political nature of many recent risks to social computing research was another major theme, already discussed in §3.1.2. This has led to researchers in this space needing to account for political harassment as a cost of doing research and to arrange for access to legal defence funds, legal budgets within grants, or insurance to mitigate risks.

Supporting professionals expressed that the risk of a researcher being sued by a big tech company was relatively low, due to potential reputational damage. But one supporting professional noted the tide may be turning:

> *« X broke the seal [in 2023]. It is still to my knowledge the only major platform that has sued researchers [over social computing research]. »*

Finally, the boundary between data journalism and social computing research can be blurry at the edges. While there are different ethical and practice norms in journalism and academic research due to their different purposes, participants observed there is much overlap in the methods and legal risks.

3.7.3 Network Analysis

Network scanning research can incur legal risks if recipients of packets are upset by research. Participants in this field pointed to existing advice on how to conduct network scans ethically as helpful [40].

Despite following the cited advice, however, one participant recounted that once, a team member accidentally made a mistake in an Internet-wide scan and caused a vendor's network devices to crash. One of the vendor's customers sent "*nastygrams*" to the participant's university: "*they emailed my department chair, the dean, the general counsel of the university, the president of the university . . . .*" The university general counsel defended the researchers, responding with a letter along the lines of, "*this is good faith security research, go away.*" Eventually, the customer stopped pursuing this threat.

Working with ISPs and setting up passive taps on network connections to facilitate research can come with legal risks under wiretapping laws (if not done carefully), an area of law that came up primarily in the context of network research. Even when participants obtained legal advice assessing a project as low risk under wiretapping law, they found that ISPs can still be unwilling to engage with research due to the high downside risk and lack of a financial benefit. Nevertheless, they reported overcoming these challenges by seeking out and working with friendly ISPs.

Measuring censorship and surveillance are common topics in network research, which we heard carry a high risk of hostile reactions from those operating censorship and surveillance systems. Participants told us about the use of anonymous or pseudonymous authorship to avoid attention (see §3.3), particularly where they live in, or have family in, the countries concerned.

Anti-censorship systems can be abused for malicious purposes, and we heard from researchers developing these systems that it can be an important precaution to avoid deliberately helping people they know are using it for malicious purposes.

Finally, participants discussed running Tor nodes, which comes with associated legal risks. As a precaution, they coordinated in advance with local IT staff to ensure that any abuse reports related to the Tor node go to the researchers rather than institutional IT staff (particularly for exit nodes). Participants also recounted using dedicated hardware that they did not mind being seized by police.

3.7.4 Cryptography and Export Controls

All three projects related to cryptography and export controls are from before 2000 when export-control restrictions on cryptography were effectively dropped. In all cases, the risk was criminal federal prosecution for violation of export control laws, which carries a significant prison sentence.

> "*I think it might've even been more, but ten years in prison was one [possible outcome], so those legal risks couldn't be ignored.*"

This issue was not observed by participants after 2000. Now, hundreds of cryptography papers are published every year without export control risks. This state of affairs was seen by participants as a success of the collective efforts of the "first crypto wars." Participants described how the risks related to cryptography before 2000 were overcome through public advocacy and changing public opinion rather than solely through legal argument. We heard that strong institutional support and support from journalists were crucial; many journalists realised this was a free-speech issue, and would affect their own profession, and so were incentivised to support researchers. Some participants wondered if community knowledge of how to stand up for the independence of researchers and intellectual freedom has survived from that time, and could be useful to push back against contemporary legal risks.

As an aside, one participant memorably highlighted the risks of "rolling your own cryptography", emphasising that cryptography engineering and building secure systems remains hard today.

> "*Any software engineer that tries to write encryption software, it's like you're a llama, so you're going to be llunch (spelled with two Ls) — that's it. And [early in my career,] I was llunch. I didn't know any better.*"

3.7.5 Malware Writing

Security researchers develop test malware for testing the effectiveness of detection tools and the feasibility of different techniques. We heard that this

type of research can raise unusual kinds of risk. For example, using a third party service for command and control of your malware, and writing about it, can cause them to issue a cease and desist, even if real malware regularly uses their infrastructure.

3.7.6 Interview Research

Research projects that involve interviewing people have their own risks (see §2 and Appendix B for our own approach to these risks). One participant discussed how they were taken by surprise by how risky the interview information they collected was, based on which they decided to revise their processes midway through a project to, for example, not use third-party or AI transcription despite the participants having consented to it, as the researchers did not feel comfortable doing so. Similarly, they felt it necessary to raise the level of the technical security measures and training of the research team to provide stronger data protection.

3.7.7 Incident Response

Security researchers dealing with operational incidents face other challenges. For example, consider a client who has been subject to a data breach through a compromise of a third party, where the breached data has been made publicly available, and the client asks for their data to be downloaded so that they can assess the damage. This involves downloading sensitive data pertaining not only to the client and their customers but many other parties. One participant recounted a situation similar to this one, where moreover, the other parties concerned had obtained a court injunction forbidding anyone from downloading the breached data. However, they added that these court injunctions have no benefit in practice as the criminals ignore them and companies do not, in their experience, enforce them against good-faith security researchers.

3.8 Ethics Review

We use the term Research Ethics Board (REB) to refer to what are in the US called Institutional Review Boards (IRBs) and in the UK often called Ethics Committees, following prior work [41, 42]. We refer to the associated review processes collectively as *institutional review*.

Research Ethics Boards and conference ethics review processes are not directly within our scope of focus, as they are designed to review specific ethical and research-process risks rather than legal risks. However, both came up frequently in our interviews as participants described them as often featuring in the process of legally risky research. Research Ethics Boards have existed for decades in the US and the UK, whereas ethics review of computer science conference submissions is only a few years old [43]; participants underscored the immaturity of conference ethics review.

Generally, participants viewed institutional review as a reasonable, necessary, and sometimes helpful process. In contrast, the prevailing view of conference ethics review was that it is a well-motivated initiative that has gone in the wrong direction, and commonly makes unnecessary, unjustified, or even unethical demands of researchers.

Background. Institutional review processes differ by jurisdiction and by institution. The canonical institutional review body at US institutions is the Institutional Review Board (IRB), whereas UK institutions have Ethics Committees. Many UK institutions delegate to departmental level for the review of most projects, only requiring central university review for the most complex or risky projects. Unlike Ethics Committees in the UK, IRBs are at the institutional (not departmental) level and the scope of IRB review is to protect the "rights and welfare of human subjects" in research—a narrower remit than reviewing ethical considerations generally—and to ensure compliance with certain regulations.[24] Appendix G provides more detailed background on how Ethics Committees work in the UK, and we refer to [41, 44, 45, 46] for background on IRBs in computer science research.

3.8.1 Institutional Review

While participants recognised that REBs are not focused on legal risks of the type our work focuses on, many reported that passing REB review can be an institutional stamp of approval that can mitigate legal risk both because it typically means the institution has been made aware of the research in advance of any legal risk situation arising, and because it may give potential adversaries the perception of institutional support.

> "*I have found that overall it's been a helpful process rather than a hindrance*"

While some researchers described REB review as more of a routine checkbox requirement, other researchers also described REBs as helpful to make sure all institutional regulations are carefully followed in every detail, and that meticulous care in doing so can mitigate legal risks in the long run.

> "*When my IRB says 'jump', I say 'how high?'*"

Conversely, participants shared that they have received harsh criticism for *not* undergoing IRB

[24] For example, research that is not human subjects research is outside the scope of IRB review, but may still involve ethical considerations.

review—most notably, from conference ethics reviewers—even when their research is clearly outside the scope of IRB review, and such review would simply have resulted in an "exempt" determination. Despite that the research was not human subjects research, participants reported that such criticism gained substantial traction and was sometimes a substantial barrier to publication, even when participants later obtained an "exempt" determination—the subsequent criticism then being that IRB review cannot be done post-hoc. In one case, a researcher reported being told by conference ethics reviewers that they should have pushed back against their IRB's "exempt" determination and insisted the IRB review the research. Participants speculated that this could be related to ethics reviewers lacking direct experience with IRB processes and/or the relevant research sub-area.

Some US-based participants told us that due to such experiences, they have adopted a policy of getting a "*defensive*" IRB "exempt" determination for basically every research project they work on. This is understandable at an individual level given the above-described experiences, but bigger picture, seems to reflect perverse incentives to waste the time of both researchers and IRBs on IRB review for projects that are out of scope.

One industry participant expressed relief that industry research is not constrained by REBs, as they see the REB process as excessive and onerous.

NIH Certificates of Confidentiality. We heard that obtaining a Certificate of Confidentiality (CoC) from the US National Institutes of Health (NIH) [47]—which is separate from IRB review, but may need to be indicated in IRB applications—can help protect research participants' data against unwanted disclosure, and provide a strong basis for institutions to help researchers protect such data against disclosure. In order to obtain a CoC, researchers need to argue that their research is health-related or within the "NIH mission" [48], which may not seem immediately relevant to most computer-science research. However, we heard about computer-science research that successfully received CoC protection, based on which participants encouraged computer-science researchers to consider whether their research might be able to be described or extended to include a health-related aspect. That said, given the US political climate (§3.1.2), another consideration when applying for a CoC is whether such an application is worth it when it might draw federal government attention, possibly heightening legal risk.

UK vs. US. Participants who had experience having their research reviewed by UK ethics committees shared that review was typically completed comparatively quickly (such as within a week or two) based on a quick write-up of the research, and was usually performed by researchers well acquainted with the relevant research sub-field and its norms. Perhaps relatedly, they often expressed that feedback received was often insightful and helpful for the research, rather than merely an exercise in compliance or box-ticking. We note that more than one of our participants were members of their UK institution's ethics committee.

In contrast, regarding IRBs in the US, participants told us that review times tended to be comparatively longer and require a more complex submission process (although process varied by institution), and that IRB members were not always familiar with the relevant research sub-field, which sometimes led to some difficulties in communication of the techniques and implications of the research. Although some participants related that their IRB's feedback genuinely helped strengthen their research, others saw IRB review as more of an exercise in compliance that did not substantively help their research.

3.8.2 Ethics Review at Conferences

Participants expressed frustration about conference ethics review processes. While acknowledging that ethics review can be a useful and appropriate part of the review of conference submissions in principle, participants generally felt that in practice, ethics reviewers often overstep their expertise and scope, and sometimes make impractical, unproductive, or even unethical demands of researchers.

Given the scope of our work, our participants were discussing their experiences of ethics review *of their legally risky research*. In this context, 4 researchers remarked on the apparent lack of familiarity of their ethics reviewers in their research subfield and its norms (most notably, vulnerability research), and also on reviewers' lack of familiarity with legally risky research, as factors that led to friction, negative reviews, and/or unproductive demands from reviewers in the ethics review process.

> "*The community expectations in the academic security community, particularly around disclosure, have gotten divorced from reality . . . I think [some ethics reviewers are] people who don't necessarily do a lot of offensive security research sort of philosophising about what they think the ideal world should be, without understanding the realities of what it's like to try to disclose to a company.*"

Some researchers reported being asked by ethics reviewers to do coordinated disclosure in ways that are inconsistent with research community norms—inconsistent not only according to participants themselves, but also according to our findings about disclosure norms (discussed in §3.7.1).

> "*[R]eviewers can . . . miss the context and point of norms. . . . Community norms can be misapplied.*"

One example involved reviewers asking a participant to do coordinated disclosure for usability studies (see also §3.7.1, discussing rationales for coordinated disclosure).

> "*The idea of responsible disclosure is to allow patches and for usability studies this doesn't make sense. But it's happened to us twice now.*"

> "*It's this contrived situation where, you know, you need to be aware of it in your future designs, maybe it feels like I taught you a new thing, but it's not like 'oh my god your product is broken.' We used to not even disclose those, because why make noise. Then [conference expectations] moved, like, if you touch hardware you need to do responsible disclosure, period, you have no judgment call anymore.*"

More than one participant reported being asked by reviewers to contact every end user affected by their findings, which in one case was impossible based on the information available.

> "*The fact is that there's been multiple reviewers who have been like, I think you should really try to disclose to all the end-user customers. But all we have is an IP address which was valid two years ago; this is not feasible.*"

Even where researchers do have end-user contact information, such an exercise would require sending mass unsoliticited communication quite possibly in contravention of applicable data protection laws.

Participants also discussed submitting research for publication where the coordinated disclosure had broken down or the recipient never replied — which, as §3.7.1 notes, appears very common.

> "*I've had a number of reviewers ask, 'why didn't you try harder to disclose to this vendor who's been ghosting for you weeks?" It's like, well, what are we supposed to do . . .*"

Another participant recalled a dispute with a reviewer about releasing exploit code. While our participants told us it is good practice to release the minimum information necessary to verify an problem and take mitigations, and to withhold certain details in order not to facilitate exploitation (§3.7.1), this reviewer disagreed.

> "*I once had a reviewer that asked that we release an unmitigated vulnerability in the name of open science. . . . I'm like, no, I'm not going to release weaponised code for an unmitigated vulnerability that can't be mitigated. That's called being irresponsible. The reviewer doubled down.*"

Participants also shared experiences of ethics committees suggesting that research should not be published because it violated "the law," and found this problematic for several reasons as follows. First, the ethics reviewers rarely specified which jurisdiction's law they were concerned with enforcing, but given the international nature of the research community, it is not coherent to enforce any particular jurisdiction's law.[25] Second, individuals on ethics committees are generally not qualified to assess or enforce the law, and that is not their stated purpose. At least sometimes, their expressed understanding of "the law" was incorrect.[26] Third, and most relevant for our present context, our field has a long history of essential public-interest research being done in legal grey areas. As such, participants expressed the conviction that we, of all fields, should be cognisant that while legality may be one consideration to weigh, it should not be determinative of suitability of of research for academic publication. We refer to Albert and Grimmelmann [49] for additional relevant discussion on the distinction between law and ethics.

Participants told us that ethics reviewers often ask them to get IRB approval for research projects that are outside the scope of IRB review.[27] Even when the authors responded by obtaining an "Exempt" determination from their IRB, the reviewers remained dissatisfied, suggesting the "Exempt" determination should have been obtained before the research started. One participant reported they have started to do IRB applications for most of their research due to this trend, even though they know it is out of scope for IRB review.

> "*[For this type of research,] going through the IRB process is literally just a hoop-jumping exercise to satisfy the reviewers because the IRB will [not review it]; they just do not do that.*"

Another participant told us that once, even with an advance "Exempt" determination, ethics reviewers

[25] Also, with double blind review, reviewers cannot know which laws the authors of a submitted paper are subject to.

[26] One participant reported receiving legal advice from a qualified lawyer that their research methods were legal; an ethics reviewer suggested they were not. When they wrote a rebuttal to this effect, the reviewer retracted their concern. However, in general, researchers should not have to hire a lawyer to rebut reviewers' speculations about illegality.

[27] We wonder whether legally risky research is particularly susceptible to this sort of request.

disagreed that the work was outside the scope of IRB review, and suggested that the authors should have pushed back and made their IRB review it. The participant concluded by remarking that not only do they believe there is no mechanism for a researcher to force IRB review, but also, "*we should not be telling the IRB how to do their job*".

That said, participants also discussed positive aspects of conference ethics review, and generally felt the community's ethical practices in its early decades were ad hoc in a way that is unacceptable now. Distinct participants told us:

> "*[In the early days,] people weren't taking these things as seriously . . . [N]owadays if you submit a paper you have to write about the ethical impact and so on. Back then it was a little more loosey goosey . . .*"

> "*I think today, [the research we did a decade ago] would be triple flagged [for] ethics committee review. . . . I would flag it to the conference ethics committee if someone did [that] now.*"

One participant suggested that some of the discontent around conference ethics review may come from community growth and diversification.

> "*The ethics complaints and . . . diverging views on how things should work, I feel like this has been changing over time and I think a lot of it may actually be due to the security community becoming more diverse, so there's a lot more Europeans and there's a lot more people from China and so I think that different cultural and legal settings are contributing to this.*"

3.9 Other Notable Themes

Our analysis concludes, in this final subsection, by addressing significant supplementary themes that were not captured in the earlier subsections: academic freedom (§3.9.1), non-disclosure agreements (§3.9.2), advocacy for change (§3.9.3), the Trusted Research programme in the UK (§3.9.4), non-legal risks (§3.9.5), and additional notable individual experiences (§3.9.6).

3.9.1 Academic Freedom

Concerns about academic freedom were a notable recurring theme in our interviews. Many researchers considered constraints on what kind of research they undertake and what they can say and publish to be unacceptable in academic contexts. Distinct participants told us:

> "*It was important to me that [no company] have editorial control over what [my] students said later, over publications. I enforced that . . .*"

Participants expressed serious concerns about the use of legal threats to stifle free expression and academic freedom.

> "*I worry that good research will be blocked, and particularly, will get blocked by powerful interests.*"Q33

One participant noted that threats designed to attack academic freedom differ in nature from hostile reactions to research from vendors.

> *« [W]e should not be mixing technical software vulnerabilities with content that is . . . disfavored by the government. »*

We heard that at least one academic institution shut down research programs and laid off researchers disfavored by powerful interests.

> "*It's an incredibly well resourced place and I would've liked to see them fight harder . . . When an institution as well respected as [this university] makes the decision to roll over like it did [and] choose[s] not to defend the right to free speech and academic work, you send the message that these attacks will work.*"Q34

But 12 other participants reported very supportive institutions who protected their academic freedom. One participant recalled checking in with their university about a planned project which could incur legal risk to both the researcher and the university, even though the participant intended to conduct the project in their personal capacity.

> "*We went to [the university] and said 'look, we're thinking of doing this . . . [but] if [you're] not comfortable with [it], we won't do it,' and [the university] said 'you should do it, right, it's about academic freedom, it's about doing the right thing, you should feel free to do what you want.'*"

See also §3.1.2 (on the US political climate), §3.5 (on institutional support), and §3.9.4 (on UK programmes that constrain research collaboration).

3.9.2 Non-Disclosure Agreements (NDAs)

Researchers commonly reported being asked to sign non-disclosure agreements, which typically had broad restrictions on researchers' ability to communicate research findings. We heard that such restrictions are often boilerplate.

> *« I think the NDAs vary, but I think Facebook's standard one required pre-publication review and had very broad definitions about the kinds of data you were not allowed to publish. »*

While signing NDAs was a common practice in many industry contexts, especially when studying a product or service to which a researcher is given special access for research, researchers in academic contexts typically avoided signing NDAs if at all possible. Distinct researchers told us:

> "*If we were trading something, sure — that's what T&Cs are typically for, that's why we agree to them. But in a world where we dump free research for a multi-billion-dollar company and then they try and gag us with T&Cs and NDAs and so on, it doesn't sit right.*"

> "*[T]he problem with NDAs in security research is that security as a research field relies entirely on transparency and sharing.*"

We heard requests to sign NDAs occurred during all sorts of interactions with companies: in contracting or employment relationships, in vulnerability disclosure interactions (within or outside the scope of bug bounty and VDP programmes) — where generally *researchers* are the ones disclosing new information to companies — as well as when simply visiting a company (e.g., to attend a research event to which the researcher was invited).

> "*At the entrance of the building, I was presented — ambushed! — with a ridiculous NDA that I had to sign. Believing that a straightforward reading of the NDA would prohibit [a research] project I had in mind, I tried not to sign it. Eventually, I signed it. . . . I was . . . deeply uncomfortable [about] the terms . . . and that I was forced to sign as a condition of attending without knowing about it ahead of time.*"

Another participant recounted a similar experience where they had been invited to give a talk at a company and were met with an NDA as a condition to enter the building, which included a condition that the researcher never say anything negative about the company. The participant found the condition unacceptable and refused to sign. They nearly gave their talk in the lobby before eventually being given a modified agreement that they found acceptable to sign. Other researchers also mentioned being asked to sign agreements containing this type of condition, and endeavouring not to sign them, with varying success. Standing their ground and repeatedly refusing worked for some.

3.9.3 Advocacy for Change

Some researcher participants had gotten involved in advocacy efforts to reduce legal risk for researchers.

> *« I think there has been a tremendous amount of progress as a result of all of our collective advocacy efforts, and in my opinion, [researchers provide] a great example of a community that is adding value. »*

Such work can take many forms. Apart from efforts to reform laws at the national level, one participant noted, *« on a policy basis, local activism and advocacy really do matter . . . »* Research and public communication raising awareness of the impacts of legal risks on research can also be an important basis for policy change (see also §6).

Participants also noted the importance of support resources for researchers, and of building alliances with stakeholders — especially those in a position to create legal risk for researchers. For example, one supporting professional noted that an employee at one of the biggest technology companies once told them that their company has never made legal threats to good-faith security researchers; the supporting professional commended them, and urged them to commit to it in writing. Such voluntary initiatives can serve as a valuable stopgap pending the slower process of more systemic reform.

3.9.4 UK Trusted Research

> "*You cannot collaborate with just anyone.*"

In the UK, there are various mechanisms intended to prevent certain kinds of research becoming available to certain kinds of people in certain countries. One of these is 'Trusted Research', a programme which began in 2021 [50] and requires universities to constrain collaborations with researchers based in certain countries or institutions (e.g., which have strong links to foreign military or intelligence services) that are deemed problematic. Since 2024, UK universities are required to abide by Trusted Research and Innovation expectations in order to be eligible for UK Research and Innovation funding (the source of most UK government research funding) [50].

Participants mentioned China, Russia, and Iran as being countries in scope, but also that it was often unclear and hard to get answers as to whether a particular country, institution, or researcher is considered problematic. The guidance talks about due diligence, protecting intellectual property, and not sharing IP with researchers who do not need it [51]. Participants shared the impression that the programme is mostly about preventing technologies

of national security importance such as quantum, AI, and some aspects of cybersecurity from being adopted by hostile states. The consequences of collaborating with the "wrong" researchers and sharing research that is not supposed to be shared were unclear to participants; however, they speculated that the consequences would be sufficiently bad that it would never be worth the risk, and would likely include university disciplinary procedures.

> "*[N]eed to make sure that I'm not going to be liable for anything that could potentially go wrong . . .*"

Researchers told us became aware of risks related to Trusted Research following internal briefings by their universities or through administrative staff advice in the context of particular projects. They reported managing these risks by involving senior colleagues in decisions about whether to proceed with collaborations, and seeking advice from more experienced researchers. We heard that these restrictions have blocked the academic implementation of algorithms from open-access published research in open-source software, undertaken by foreign nationals from disfavored countries who were to be supervised by UK academics.

> "*I'm now in the very difficult position where I would love to work with these people, their published work is good, they cite my work, they like my work, so it makes sense for us to work at the academic level, but because of the rules and the way that it's still untested waters as to how these restrictions will be applied and enforced[, I have not been doing so.]*"

Another similar mechanism is the Academic Technology Approval Scheme (ATAS), which requires certificates from the Foreign, Commonwealth & Development Office (FCDO) for students to study certain subjects "in specific sensitive technology-related fields" at a Masters or PhD level [52]. PhD students must obtain this certificate prior to getting a visa to study in the UK. In the application form they have to detail what they will be studying, prior employment, prior publications, and referees. According to participants, ATAS applications are regularly refused for reasons not disclosed to potential supervisors. One participant who worked in a sensitive field said they did not find it worth considering applicants from certain countries due to the likelihood they would not get an ATAS certificate.

> "*[I]t's not worthwhile trying to actively recruit from overseas because of the risk of ATAS not being granted.*"

They added thats this significantly hampered their ability to get PhD students, as most students applying with funding were from countries where they would be unlikely to get an ATAS certificate.

Participants raised concerns about academic freedom related to these UK initiatives, related to increased institutional review of what kinds of research and collaboration are taking place, the power of university authorities to forbid research from taking place, and the lack of clarity around who is making the decision to block certain collaborations and how that interacts with the normal university chain of command. See Q33.

Another related theme was the lack of clarity on which topics and countries were disfavored, and the broad chilling effects resulting from such ambiguity.

Similar issues also arise in other countries. For example, eligibility for some European Union Horizon grant calls is restricted due to national security concerns [53], and US federal grant funding for research has become increasingly politicised in recent years [54, 55, 56].

3.9.5 Non-Legal Risks

> « *[Y]ou know, we've talked about legal risks, but there are other risks too — harassment is a big one. A lot of researchers have been targeted through harassment campaigns online, especially women researchers.* » Q35

There are also important *non-legal* risks that researchers face. While these are largely outside our scope of focus, they sometimes correlate with legal risks. Many participants discussed being personal and professional attacks in the press, impugning their motives and character or belittling them, as well as harassment. As Q35 and Q14 highlight, these may depend on researchers' demographics.

One researcher told us they turned to university lawyers for help after receiving death threats, and were sorely disappointed by the response.

> "*I got no advice from the university lawyers whatsoever. They either ignored me or told me it was my own problem for having gotten death threats.* "

There was one instance where a participant reported being threatened with physical violence by a government agent, in the context of an interaction related to criminal research-related legal risk. It went no further than a verbal threat, and the participant stressed that they strongly believed this to be a rogue agent without institutional backing.

Moreover, many participants discussed non-legal ethical considerations and challenges that they face in their research, and concern about "doing the right thing" when faced with ethical challenges.

See also §6 for prior work on non-legal risks.

3.9.6 Other Concerning Experiences

We briefly highlight three notable participant experiences not covered elsewhere in the paper.

- We heard about two incidents where researchers carrying physical research equipment in public places had the potential to raise alarm from law enforcement. In one such case, the researcher in question was arrested. Describing the other case, a participant recalled:

  > "*[W]e just realized that we had gone through the [public transit system] with a bag full of wires and stuff, so then when we went back home we got a taxi . . . So, I don't know if that counts as a legal risk, being shot by the police, but yeah, that was a risk that we thought about . . .*"

- A participant who was imprisoned recounted writing a letter to a government official who had publicly discussed the importance of protecting researchers from legal risks:

  > "*The jail intercepted the letter and told me that I am not allowed to write [to them] and if I ever do that again they're going to throw me in . . . solitary [confinement].*"

- A US-based participant recalled being taken without explanation to an interrogation room and questioned about their research when reentering the US through a major airport. When the participant asked to talk to their lawyer, the interrogator refused and insisted that the researcher would stay in the room as long as it took until all their questions were answered. Though seriously intimidated, through standing their ground, refusing to answer, and insisting on seeing their lawyer, the researcher was eventually released.

4 Recommendations

4.1 Researchers

While we find legal risk experiences are highly contextual, our analysis points to three broad recommendations that researchers navigating risk can unilaterally put into action:

1. *think through legal risks early*;
2. consider the *role of stakeholder incentives and public communication* in how legal risk situations play out; and
3. build a *support network of trusted colleagues* and lawyers to turn to in legal risk situations.

Related to (1), discussion of legal risks early on within research teams serves to establish a shared understanding of and approach to legal risks that works for everyone. Advance planning can facilitate quick responses and reasoning under the stress and time pressure that may accompany legal risk incidents, as well as enable proactive precautions throughout their research process. Early-stage planning could benefit from a brief attorney consultation, for those who can afford it or access *pro bono* counsel (see Appendix F for an example).

On (2), the relevant stakeholders vary widely depending on context, including those within an institution, entities instigating the risk (if any), and others such as regulators, the media, and more. Any of these might have interests at odds with researchers', and anticipation of those incentives can help risk managements. Especially when faced with PR attacks, researchers may benefit from consulting PR experts. See also Q18.

> « *[Researchers] need to understand that a big part of the strategy on the other side is a PR strategy. . . . Researchers and their supporters need to counter that narrative, to make sure the public better understands what they're doing, why they're doing it, why it's important, and why the threats are threats not just to the research, but to the public . . . Researchers are engaging more and more in this way, but it's not how most researchers are trained to engage with the public . . .* »Q36

On (3), advisors, mentors, and senior colleagues were highlighted by participants as helpful, as well as those with specific relevant experience: e.g., in a research area, with an institution, with certain kinds of threats, and so on. A common strategy to connect with lawyers was at conferences and through existing contacts who are lawyers (including law school colleagues; but not all law schools have this specialised expertise).

Our detailed *strategies*, *precautions*, and *advice*, covered at length in §3.3, discuss many more concrete, actionable steps towards our recommendations. Specific strategies and precautions' relevance will depend heavily on researchers' context.

When To Get a Lawyer. We hope researchers leverage our participants' diversity of experiences with lawyers, detailed in §3.4, to inform their own approaches. Our analysis suggests that legal advice can be valuable for any kind of legally risky research at any stage, perhaps especially when researchers are unsure if an activity is risky, when their institution's interests might not be aligned with their

own, or when facing litigation risk. Participants recommended contacting lawyers early if possible, to avoid avoidable practices that heighten risk (e.g., Appendix F). But see Q20.

Watch Your Stockpiles. Given the prevalence of stockpiling vulnerabilities (§3.2.3) and the risks this practice poses for researchers, digital security and privacy, and national security (see §5), we encourage researchers to reflect anew on their stockpiling practices, to limit and closely monitor stockpiled information, and to reinforce operational security.

Tell Our Stories. We encourage researchers to speak out if possible, to raise awareness and argue for change (§4.2).

4.2 Other Stakeholders

Supporting Professionals. Our three key takeaways regarding supporting professionals are: (1) proactive outreach to researchers via conferences, forums, and other networking importantly raises general awareness of support resources; (2) proactive outreach to individuals known to be facing legal-risk situations can have great impact; and (3) for attorneys specifically, initial researcher scepticism about lawyers is common, so investing effort in rapport-building can be essential.

Academic Institutions. In light of participants' varied experiences with institutional support (§3.4, §3.5), we encourage institutions to better support researchers and to defend academic freedom against mounting attacks, and especially not to "obey in advance" (see also Q5, [57]). That said, the reality is that institutions' and researchers' interests may not always align perfectly, and what researchers really need may be support from outside their institution (§3.5, Q27). We hope that raising institutions' awareness of the broader legal-risk landscape and impacts of research and researchers (which we document) will help institutions recognise that researchers' expectations may be misaligned with what they can provide, to communicate clearly to researchers the extent of their ability to help, and to support researchers in finding external support resources where institutional support is not enough.

Conferences & Reviewers. In light of participants' varied experiences with reviewer reactions to legally risky research, we encourage conferences to develop explicit policies that acknowledge the difference between law and ethics and the importance of certain kinds of legally risky research, and state that while legality is an important consideration, the conference's role is not to enforce the law of any jurisdiction. Recognising that doing so may itself entail legal risk, we recommend conferences consult an attorney with expertise in this area, in crafting such a policy. Background reading about the history of legally risky research and the law/ethics distinction (e.g., [49]) may be helpful to get reviewers from diverse contexts on the same page. Moreover, our findings suggest that sub-field expertise and norms can be critical to fairly reviewing legally risky work; we encourage conference ethics committees to seek input from sub-field experts.

Companies. Companies can help by not legally threatening good-faith researchers, and not suggesting that research activities are criminal or otherwise legally actionable.[28] Our findings suggest that both companies and researchers could benefit from companies setting up dedicated contact methods for external and internal research findings, that include an explicit process for routing research reports to technical teams with relevant specialised expertise; moreover, we encourage companies to offer contact methods not tied to contractual conditions. We recognise not all companies may have the resources to do this, and recommend leveraging support and expertise from existing organisations like disclose.io, HackerOne, or CyberUp (see Appendix A). We encourage companies to embrace programmes for researcher data access without imposing contractual conditions that hinder open science (see Appendix A). Finally, we highlight a participant's suggestion that voluntary pledges from companies not to threaten good-faith researchers could help mitigate legal risks in the short term pending broader systemic change.

Law Enforcement. We encourage law enforcement to exercise prosecutorial discretion in favour of good-faith researchers (and thereby, national security; see §5), to establish internal processes by which "hacking" cases are reviewed by experts aware of the legal-risk landscape for researchers before a case is pursued, and to put into writing clear formal or informal policies not to pursue cases against good-faith researchers.[29] Law enforcement may also play an important role in emphasising the national security implications of the legal-risk landscape (see, e.g., [60] on encryption).

Government Agencies & Intermediaries. Government agencies, especially related to cybersecurity[30] or particular industries, can critically reduce legal risk for researchers by: raising public awareness of the importance of research, especially within industry and policy circles; by promoting industry norms and standards in favour of accepting researcher feedback; by coordinating the dissemination of research (possibly as a disclosure intermedi-

[28] There are numerous past examples of such suggestions by companies. See, e.g., [58, 59].

[29] The 2022 DOJ charging policy [29] and DMCA exemptions are good examples; participants indicated that even more clarity and enhanced scope of protection would be even more helpful.

[30] E.g., CISA in the US or NCSC in the UK.

ary while keeping researchers anonymous (§3.7.1)); and by raising researchers' awareness of such coordination possibilities. We encourage agencies and other potential intermediaries to maintain their existing support for researchers and consider expanding their scope.

5 Discussion

Who Will Do The Research?. Researchers with the privilege of working for well resourced and supportive institutions and of not being in vulnerable groups (§3.2.6) may be especially needed to speak out, to support colleagues and students, and to undertake research others might not risk doing.

National Security. Cybersecurity has become a pressing, pervasive problem at a global scale as many critical systems now rely heavily on digital technology — not only for digital economies, but also for national security. In both the UK and the US, authorities have acknowledged the urgency of the issue, as well as the national security concerns around the growing cybersecurity capabilities of actors they see as adversarial [61, 61, 62, 63]. The UK's National Cyber Security Centre (NCSC), for example, reports that in the past two years, China's offensive cybersecurity capacity in both skill and volume has massively increased, partly driven by buying access and vulnerabilities from private companies. Against this backdrop, our findings — indicating that the solicitation of the sale of security vulnerabilities from security researchers has long been widespread and normalised, and may have accelerated in recent years (§3.2.3) — are especially concerning. Ensuring that researchers of all stripes — not limited to "good-faith" researchers — feel safe disclosing vulnerabilities to impacted parties in a coordinated manner is therefore crucial to national security, as the risks that otherwise the vulnerability will accidentally or deliberately fall into the hands of adversarial nation states are growing rapidly. Ensuring the same is moreover essential to the upkeep of a safe and secure Internet and digital technologies at a global scale, protecting those who use and rely on these technologies from their own and other governments as well as other adversarial actors, wherever they are located.

> "*Adversaries who are actually adversarial will not say, 'oh, well it says in the terms of service we can't do this, and therefore we won't do it.'*" Q37

The US National Security Agency has acknowledged the importance of this issue in an initiative to establish coordinated disclosure programmes to work with researchers, a collaboration with CISA as well as cybersecurity agencies in Japan and the Netherlands [64]. CISA and NIST have a related programme and recommendation, respectively, for federal agencies to share and receive cybersecurity information from researchers [65, 66].

A participant noted that there may be a parallel in today's situation with the evolution of past national security arguments opposing encryption.

> "*The reasons that they had [for restricting cryptography] in the years following World War II were good reasons at the time, but with the arrival of personal computers and the Internet, those reasons were not as important as protecting people's privacy and business privacy you know and also even national security communication . . . In the '90s, the government used a national security argument against the spread of strong cryptography, but today, the national security argument favours widespread use of strong cryptography.* "

We are not the first to point out the national security concerns related to overbroad anti-hacking laws (see Pfefferkorn [67]). Our findings exacerbate the concerns already articulated by others.

One of the implications of this backdrop is to the Vulnerabilities Equities Process [68], where the relative benefit of disclosing government-discovered (or intermediated) vulnerabilities to affected parties is weighed against the benefit to the government of exploiting those vulnerabilities. A higher risk of discovery and exploitation by adversaries counsels a shift in the balance towards disclosing vulnerabilities so that they can be fixed.

AI-powered Vulnerability Finding. Following our interviews, Anthropic's release of Claude Mythos made waves in the security community. Mythos is a specialised tool with cybersecurity and vulnerability finding capabilities; while automated vulnerability discovery was possible before[31] Mythos and its competitor tools have made vulnerability discovery significantly faster and more scalable in recent months [70, 71, 72, 73]. This acceleration has escalated the urgency of remediation to maintain the historical balance between attack and defense, as well as the need for coordinated industry responses [74, 75, 76] — exacerbating the risks associated with hostile attitudes towards security research findings which we find are still common (§3.7.1). For example, in May 2026, Anthropic reported that they had discovered 1,596 vulnerabilities across 281 open-source projects and to their knowledge at the time of writing, 97 of those had been patched [77]. These advancements heighten the critical security and national security risks discussed above.

[31] Google's Big Sleep first discovered a zero day in 2024 [69].

Legal Reform To Date. Both the UK CMA and the US CFAA have been amended since their initial passage in 1986 and 1990, respectively. The amendments have generally expanded the scope of liability under the two laws, increasing risk for researchers over time [78, 79]. Proposals to narrow the scope of liability to protect research activities have been made, but to date have not been enacted [12, 80, 81]. In late 2025, the UK government committed to revise the CMA to better protect researchers [82].

In contrast, the US Copyright Office holds a triennial rule-making procedure which establishes temporary exemptions to liability under the DMCA, which has successfully strengthened protections for researchers over the years, albeit as temporary exemptions rather than permanent reform [83, 84, 85].

In the absence of statutory reform to reduce risks for researchers, participants noted that case law and prosecutorial discretion have also been helpful in some cases. The DOJ has issued non-binding guidance indicating that good-faith security researchers should not be prosecuted [29]. Distinct participants remarked:

> « *The actual net legal risk has gone down, in the formal sense of what the law says about what's happening. There's still huge areas where we literally have no case law* »

> « *[P]rosecutors have realized that this is having a chilling effect on research and have taken, you know, I don't think they're sufficient steps, but have taken some steps to try and reduce that threat* »

From a cybersecurity perspective, encouraging the disclosure of discovered vulnerabilities is in the public interest even for bad-faith researchers. Other public-interest considerations suggest that bad-faith researchers should face deterrents and/or consequences for harm they cause. Balancing these interests is a delicate challenge that is an important open area for future research and policymaking.

Finally, we share a participant's reflection taking a broader view of legal reform beyond research:

> "*I'm just an ordinary person out there now today. Most people aren't that lucky, even if they were wrongfully convicted. Most people stay wrongfully convicted, and that's just the way that the . . . justice system is. I think there needs to be reforms, not just in cyber policy, which is atrocious . . . but the criminal justice system as a whole needs to be reformed.*"

Advocacy Impacts. Advocacy by researchers and supporting professionals has been impactful (see also §3.9). Some notable examples include: filing amicus briefs[32] in a landmark 2021 US Supreme Court case, *Van Buren*,[33] which narrowed the CFAA's scope, reducing risk for researchers [86]; and submitting comments to the US Copyright Office's triennial rule-making proceedings [87, 84], including empirical evidence on legal risks [4], resulting in explicit exemptions to DMCA liability for good-faith research which have strengthened over time; organising convenings and comment submissions which led to revisions to the NIST Cybersecurity Framework in support of vulnerability disclosure processes [88]; and participating in the CyberUp Coalition [89], a UK advocacy organisation (see Appendix A), whose efforts contributed to the UK government's recent commitment to revise the CMA [82]. This list is not exhaustive. We encourage readers to consider getting involved.

6 Related Work

This section overviews related academic work. Appendix A lists other useful resources for researchers facing legal risk.

Empirical Literature. Hall and Adams' 2018 whitepaper analyses how security researchers decide which projects to pursue, in the context of legal risk, including an interview study with 20 participants [90].[34] They focus on this narrower question, whereas we study researchers' experiences at larger scale and not limited to project choice (e.g., broader personal and professional impacts, experiences with lawyers, institutions, and more). Akgul et al. survey and interview bug hunters on the challenges and benefits of bug bounty programmes and their motivations [91]; they focus on vulnerability research and bug bounties specifically, whereas we consider broader research, experiences, and impacts. Gamero-Garrido et al. offer an empirical study of the DMCA's chilling effects on vulnerability research [4], and Wang et al. conduct an empirical study showing chilling effects of computer crime law enforcement on hacker forums [5]; both focus more narrowly on chilling effects than our work. Hrle et al. examine Redditors' discussions and advice to each other on legal risks of port scanning [92].

Legal Literature. Park and Albert's 2024 guide examines US-based legal risks for security re-

[32] These are court filings from parties not involved in a case, weighing in on broader implications of the case beyond the parties involved.

[33] *Van Buren v. United States*, 593 U.S. 374 (2021).

[34] They had two female participants and noted a more gender-representative study would be important future workespecially given the field's gender imbalance; as discussed in §2, we had a better gender balance than [90], roughly mirroring that of the field.

searchers [6], Xiao et al. expose risks to researchers under anti-fraud laws [14], Sellars studies risks of scraping under the CFAA [10], and Burstein offers an earlier overview of legal risks with an emphasis on network research [93]. Kumar et al. and Evtimov et al. studied legal risks for adversarial ML research [11, 94], Soghoian and Ohm studied legal risks for phishing and network research respectively [7, 95]. Various legal scholars have written about the need for reform to protect research (e.g., [96, 97, 9]).

Non-Legal Risks. Doerfler et al. interviewed researchers who have experienced harassment [98]. Marwick et al. provide a guide to conducting risky research safely, considering risk not limited to legal risk [99]. Gilbert, Matias, and Fiers interviewed 14 independent researchers to investigate the risks and threats they faced, providing advice on how policymakers, research ecosystem supporters, and researchers can mange risks [100]. Abdo et al. offer a framework for researcher access to public platform data [39]. Thomas et al. study ethical issues in research using datasets of illicit origin [42].

Vulnerability Disclosure. Concurrent work by Ramulu et al. presents a qualitative interview study on vulnerability disclosure experiences [101], which is complementary to our work: we focus on legal risk experiences, including vulnerability research as one among many types of risky work; they focus in depth on disclosure, documenting legal risk as one among many aspects. There is an extensive literature on vulnerability reporting and the importance of attack research, e.g., [102, 35, 103]. Kohno et al. provide a set of scenarios for considering ethical dilemmas in vulnerability disclosure [104]. Elazari has spoken and written about the effectiveness of bug bounty programmes from a legal perspective [105, 106].

7 Conclusion

We have presented an unprecedented qualitative mapping of the broad landscape of legal risks in computer science research, accompanied by actionable recommendations, strategies, and advice. We have shown that researchers have a wealth of compelling untold narratives from *« sticking their heads out above the parapets »* to do essential work to keep systems we rely on safe and accountable.

> *« [R]esearch is done for the public benefit[. S]ecurity researchers play a really important role in assuring the security and privacy of technology, and have done historically. That's not something I think the general public understands fully. You don't see concrete stories of this that often — of the civic mindedness of researchers. So few stories of civic mindedness in general these days. »*

We hope our work, in telling some of these stories, both helps researchers in the immediate term, and more broadly raises awareness and provides a basis for change.

Acknowledgements

We thank our participants for their interviews and for sharing experiences with us that were often sensitive, difficult, or stressful. We enjoyed interviewing you, and appreciate your time and candour. We also thank those who suggested participants; we are not naming you, to enhance the anonymity of our participants. Finally, we are grateful to: Belen Pisaniello for excellent research assistance throughout the project; Rosanna Bellini, Ben Collier, and Miuyin Yong Wong for advice on and review of research methods; Adam Sealfon for advice on aggregating statistics; Nicholas Spooner for interview practice; Stephanie Chen and Kaitlin Ponder for contributing background legal research; and our anonymous USENIX '26 reviewers and shepherd.

A Useful Resources for Researchers Facing Risk

We welcome suggestions of additional resources. We plan to maintain a repository of useful resources at [107] and keep it updated with community suggestions; please check [107] for an up to date list.

Legal Advice. In the US, legal advice and support is sometimes available from the Electronic Frontier Foundation (EFF) [108] (particularly for security research) and the Knight First Amendment Institute [109] (particularly for social computing research).

The only law firms we are aware of specialising in in research-related risks are Albert Sellars [110] in the US and pwn.legal [111] in the UK. They both sometimes do *pro bono* work. Additionally, there are other individual lawyers with relevant expertise.

A few universities in the US have legal clinics that may be able to support students or faculty facing legal risk; examples include the BU/MIT Student Innovations Law Clinic (SILC) for Boston University and MIT students [112], the Samuelson-Glushko Technology Law & Policy Clinic at the University of Colorado Boulder [113], the Cyberlaw Clinic at Harvard Law School [114], and the Technology Law and Policy Clinic at New York University (at which one of the authors does *pro bono* legal work) [115].

We were unable to identify similar organisations in the UK (e.g., the Open Rights Group (ORG), which fulfils a similar role to the EFF, does not

provide legal advice). We are not aware of any UK universities with similar clinics. (See also §3.6 for jurisdictional context.)

Community Support Organisations. We heard about a number of researchers' community support organisations, some of which also engage in advocacy work. Examples include disclose.io [116], I am the Cavalry [117], Austin Hackers Anonymous (also known as AHA!) [118], a regional group based in Austin, Texas.

Legal Defence Funds. There are legal defence funds to which researchers can apply for funds to help cover costs of legal support: the Security Research Legal Defence Fund [32] and the Internet and Platform Researcher First Amendment Support Fund [33]. Even if these legal defence funds do not award funding to a particular researcher's case, they may be able to provide helpful resources or contacts. See §3.4.3 for additional discussion.

Documented Incidents. threats.disclose.io documents legal threats to security researchers globally based on research community contributions [119].

Advice for Resaerchers. Albert (of Albert Sellars, above) spoke about legal restrictions in bug bounty programmes at USENIX Security (Enigma) 2025; the video is available online [37]. Barclay collated advice on how to handle cease-and-desist letters from big technology companies [120]. The Research Support Consortium provides tips for scholars engaged in public communication about elections [121]. See also our Related Work section (§6), including Park and Albert's guide to legal risks of security research under US law [6], and other similar guides focused on specific areas of law. Please note that not all of the cited works are recent and they may not reflect the current state of the legal landscape.

Advocacy Organisations. There are a number of groups involved in advocacy to improve legal risks for researchers. Examples include CyberUp [89] and the Open Rights Group (ORG) [122] in the UK and the Electronic Frontier Foundation (EFF) [108], the Knight First Amendment Institute [109] in the US, as well as the international Coalition for Independent Technology Research [123]. Interested researchers may wish to consider getting involved in and/or supporting the work of such groups.

Data Sharing. Social computing research requires access to data. Recent efforts towards improving data access for research include the EU Digital Services Act, which requires certain companies to share certain platform data with researchers [124] and Abdo, Drobnic Holan, Silverman, Hickey, Ardnaudo, Merrill, Lukito, Hendrix, Dowling, Edelson, Bode, Seiling, Farid Johnson, Shiffman, Persily, Chapman, Faust, Tromble, and Ventura's roadmap for expanding access to public platform data [39].

Other Resources. The researchers' amicus brief in the 2021 Supreme Court case *Van Buren* [86] and researchers' filing to the US Copyright Office about legal impediments to security research [125] may also be of interest.

B Additional Ethical Considerations

The stakeholders in this research are: the research team; the researcher participants; their respective research teams, institutions (past and present), and lawyers; the supporting professional participants; the organisations that provided the products or services that were the subject of research that participants described ("providers"); the users of those products or services; UK and US government officials and agencies including law enforcement and those with responsibility for security; the security and social computing research communities more widely; and wider society.

Our researcher participants had already been subject to risk of harm from many of the other stakeholders, and interviewing them risked materialising potential harms and creating risk of new harms. Our research team faced risks from attracting the attention of stakeholders who might wish to harm anyone associated with particular interview participants. Anonymisation and minimisation of records helps avoid attracting unwanted attention and so protects the research team and participants.

There would be a risk to providers and users if participants disclosed to us vulnerabilities or issues not yet known to the public and that information was then compromised. Participants did sometimes disclose high-level information about vulnerabilities not in the public domain (e.g., there is a vulnerability in particular a system of a particular severity) — either in cases where the disclosure process was ongoing, or they had decided never to disclose (see §3.2.3). We did not solicit and did not receive detailed information about these vulnerabilities that would enable them to be exploited; if we had, we would have redacted that information before transcription so that we kept no written records of it. We also never learned a level of detail of data that would enable us to independently report the vulnerability. We redacted the details of which systems are impacted in our records. When participants asked us to not to publish certain vulnerability-related information within an embargo period, we respected their requests.

The disclosure that vulnerability researchers are stockpiling vulnerabilities that they do not feel they can safely disclose (see §3.2.3), and possibly selling

them to mysterious parties, presents a risk to individuals in these communities and to the public. It is a 'vulnerability' which bad actors are already actively exploiting, but disclosing it could make more bad actors aware of it. We have coordinated with community members on appropriate disclosure to affected communities.

The aim of this research is to enable researchers to minimise risk in future through a better understanding of what problems occur in what situations and how to mitigate them, and to evidence the need for policy changes that minimise risk (benefiting the security and social computing community). In turn, we hope this will increase the amount of public interest security and social computing research conducted, benefiting wider society. On a utilitarian basis, we believe this benefit outweighs potential risks to other stakeholders; on a deontological basis, we obtained informed consent from participants, who face the greatest risk, and should have the autonomy to tell (or not tell) their stories. Law enforcement and government agencies with responsibility for security might be interested both in our results and in heightening legal risk for some participants, so we took steps to protect participants from these stakeholders (minimising harm) while ensuring the results are accessible to them (maximising benefit).

Finally, our measures to protect participant anonymity went beyond standard practices, given the above risks. We insisted on anonymity for all participants, including those happy to be identified, to enhance the anonymity set for all participants. We requested explicit approval from participants for every direct quote in the paper, and allowed them to edit or veto each quote. While we made a best effort to anonymise quotes before sending them to participants for approval, participants are in the best position to know what details or even styles of expression may be identifying. Even participants who were happy to be identified were proactive in giving us tips to better anonymise their quotes when the anonymity set rationale was explained. Some details of how corner cases were coded are elided as they would be identifying.

C Glossary of Acronyms

Acronyms are generally introduced the first time they are used. We additionally provide this glossary for easy reference.

- BBP: bug bounty programme
- E2EE: end-to-end encrypted (or end-to-end encryption)
- EULA: end-user license agreement
- CERT: Computer Emergency Response Team
- CFAA: US Computer Fraud and Abuse Act
- CISA: Cybersecurity and Infrastructure Security Agency
- CMA: UK Computer Misuse Act
- CSAM: child sexual abuse material
- DOJ: US Department of Justice
- DMCA: US Digital Millennium Copyright Act
- EFF: Electronic Frontier Foundation
- IoT: Internet of Things
- IRB: Institutional Review Board
- NCSC: UK National Cyber Security Centre
- NDA: non-disclosure agreement
- NIH: US National Institutes of Health
- PR: public relations
- TOS: terms of service
- VDP: vulnerability disclosure programme

D Index of Referenced Quotes

Q1 . . . . . p 3
Q2 . . . . . p 4
Q3 . . . . . p 10
Q4 . . . . . p 11
Q5 . . . . . p 12
Q6 . . . . . p 12
Q7 . . . . . p 16
Q8 . . . . . p 16
Q9 . . . . . p 17
Q10 . . . . . p 17
Q11 . . . . . p 17
Q12 . . . . . p 17, 18
Q13 . . . . . p 18
Q14 . . . . . p 18
Q15 . . . . . p 19
Q16 . . . . . p 21
Q17 . . . . . p 21
Q18 . . . . . p 21
Q19 . . . . . p 25
Q20 . . . . . p 26
Q21 . . . . . p 27
Q22 . . . . . p 28
Q23 . . . . . p 30
Q24 . . . . . p 30
Q25 . . . . . p 31
Q26 . . . . . p 32
Q27 . . . . . p 33
Q28 . . . . . p 36

Role	Count
CEO	10
Consulting (while academic)	1
Government	1
High school	2
Independent	9
Industry	18
Non-Tenure-Track	15
Nonprofit	1
PhD Student	16
Tenure-Track	12
Tenured	42
Unanswered	1
Undergraduate	1

Table 5: Researchers' role at the time of the project

Jurisdiction	Count
Both	1
Both & Other	2
Europe	4
Other	2
UK	35
UK & Other	1
US	84

Table 6: Researchers' jurisdiction at the time of the project. 'Europe' here does not include the UK.

Q29 . . . . . p 39
Q30 . . . . . p 39
Q31 . . . . . p 40
Q32 . . . . . p 41
Q33 . . . . . p 48
Q34 . . . . . p 48
Q35 . . . . . p 50
Q36 . . . . . p 51
Q37 . . . . . p 53

E More Details On Participants and Methodology

Table 5 lists the roles of researchers at the time the project was carried out. Most projects were carried out in an academic, industry, or independent setting, but there were some more unusual settings such as academics consulting for industry, work in government, or at a non-profit. Postdocs are mapped as 'Non-Tenure-Track', established UK faculty are mapped to 'Tenured' as the closest equivalent.

Table 6 shows the jurisdiction of the researchers at the time of the project, due to our scope it is dominated by US and UK, but sometimes researchers discussed projects carried out in other jurisdictions, or which spanned multiple jurisdictions.

E.1 Details on Coding Method

Type of Legal Risk. We classified the type of potential legal risk of each project as criminal, civil, both, or not applicable, based on a rough classification of the type of research, activity, or incident described — *not* based on a thorough legal analysis of the details of the project. (Indeed, we generally did not have enough detail about the projects to do such a thorough legal analysis; and collecting such detail was beyond the scope of our project and would have created unnecessary confidentiality risks.)

We performed this rough classification primarily based on the reference table in [6, Appendix A], which lists types of research activity and the kinds of legal risk to look out for, for each. However, we adopted the following policies in addition to using [6]'s reference table:

- UK-based projects of the type that [6] suggested could be subject to CFAA liability were classified as having potential criminal liability only, because the CFAA includes both civil and criminal liability whereas the CMA has only criminal liability (see also 3.6).

- Although [6] suggests that CFAA liability for TOS violations is unlikely after *Van Buren* [126], we classified projects involving potential TOS violations as having potential CFAA liability, because many of the projects we heard about were from before *Van Buren* (i.e., before 2021). For simplicity, and also in acknowledgment of the fact that the precise scope of the *Van Buren* ruling is still not perfectly clear, we chose to classify projects consistently according to the preceding policy rather than treating post-2021 projects differently.

- [6] does not cover data protection law. We classified projects involving the collection or processing of personal data as having civil risk under data protection laws.

- Where a participant specifically told us about legal risks under particular laws and whether they were criminal or civil in nature (regardless of jurisdiction), we used this information for our classification.

The above methodology means that:

- where projects were in a legal gray area, our classification included the potential gray-area risk even though the likely treatment of that type of research is in fact uncertain; and

- types of potential legal liability that are not covered in [6], other than data protection law, are not considered in our classification (e.g., potential risks related to handling of classified materials, or under financial and banking laws, would not be taken into account), except where specifically noted by participants (as noted in the final bullet of the previous bulleted list);

- our classification of the type of potential legal risk is largely based on the US context (as [6] focuses exclusively on the US context), although our method takes certain aspects of UK law into account.

Type of Legal Threat. We classified each legal *threat* as criminal, civil, or both. Here, the classification was based on whom the threat came from and the content of the threat: if a company sent a cease-and-desist letter suggesting a researcher violated the CFAA, then we would classify that as a civil legal threat:

- even though the CFAA can also create criminal liability, because the threat came from a private party; and

- even if we would not have classified the type of potential risk as including CFAA liability, and regardless of whether the researcher indeed violated the CFAA.

If a warning or threat came from a law enforcement, that would be criminal legal threat. If a researcher experienced both criminal and civil legal threats on the same project, we classified it as "both."

Legal Representation. We marked a project as having legal representation if the corresponding researcher participant formally consulted a lawyer for advice on that project, however briefly. We do not count as legal representation any other advice of a legal nature that the participant may have received, from lawyers that were not their own (e.g., university or company lawyers).

More Detail on What Constitutes a Project. 'Project' encompasses both traditional research projects that were systematically planned, accidental discoveries, and incidents of legal risk related to the kind of research the researcher did, but that did not relate to an ongoing project. We only count projects where the participant was directly involved with the project, not their recounting of stories that happened to other people, except where it was the independent projects of students taking their class, where they advised the students. Sometimes one project is recorded for multiple experiences that the a researcher talked about as a bundled experience.

We interviewed participants with experiences that are in scope of the research, but a few of them mentioned projects that are out of scope, because of being based in jurisdictions other than the UK and the US. We included these projects in the analysis, as they form part of the experience that in-scope researchers shared with us. For such projects, we classified the type of potential legal risk using the same heuristic as already described for the US/UK, for consistency, although this may be an inaccurate reflection of legal risks in third countries.

We only coded projects for researcher participants. Supporting professionals mentioned projects they helped researchers with, but we did not code these as projects.

Project Publication. Where projects were ongoing at the time of interview, or where it was an incident of a type that would not typically result in a publication, then the publication status was coded as Not Applicable.

F Example: An Ounce of Prevention

This is a example suggested by one of our participants, with some details modified for anonymity. It illustrates how a brief consultation with a lawyer might impact a researcher's process to reduce legal risk, while preserving all the researcher's goals around the technical scope of the work.

Car insurance companies sometimes offer discounted rates if insurance clients fit their cars with certain tracking devices, but it can be unclear what exactly these devices do and what data they collect. Imagine a researcher Ria who wants to understand how these devices work. She is interested in obtaining some devices and reverse-engineering them. At the very start of the project, she briefly consults a lawyer Leo on any precautions she could build into her research design to minimize legal risk. Leo advises Ria that her legal risk will be lowest if she can avoid getting into any contractual relationships with insurers or device manufacturers. Therefore, he advises her not to obtain the devices through the usual channels (by purchasing insurance in your own name and fitting the device to your own car). Leo additionally notes that obtaining a device from an insurer under a false identity or by falsely representing you have a car to insure would be fraud, heightening legal risk, so that should be avoided as well. Ria and Leo sat down together to figure out how to obtain devices for reverse engineering without incurring any of the above risks, and eventually realized that many of these devices get stuck in cars and then forgotten about. They worked out a plan for Ria to go to junkyards and ask them to check for the device before stripping down the car, and to offer some amount of money for each device the junkyard pulled out for the research team

— thus avoiding any contractual risk whatsoever. Ria's team then proceeds with reverse engineering these devices just as they normally would.

As one lawyer participant put it, even a brief consultation with an attorney can bring to light *« things like that where . . . you wouldn't think about it you weren't a lawyer obsessed with these kinds of weird things »* and enable adjustments that can allow the research to proceed with lower risk.

G Background on UK Ethics Committees

Many UK institutions delegate to departmental level for the review of most projects, only requiring central university review for the most complex or risky projects. Departments can have quite lightweight and fast processes, such as targetting a 5 working days from submission turnaround (facilitating the review of hundreds of undergraduate/masters student projects each year in any one computer science department), while university level review generally has more process and reviews according to a fixed committee meeting schedule (e.g. monthly). Outwith medical and human tissue research institutional ethics review in the UK is not driven by specific legal requirements, but by guidelines and recommendations from funders that will only fund institutions with appropriate policies and processes in place [127, 128, 129, 130].

This gives institutions a lot of flexibility in how they implement the general principles and while they must have review processes that cover human participant research (as IRBs cover in the US) they can cover anything else they deem important [131]. UK ethics committees will often consider data protection requirements (e.g., GDPR) [132, 133] but normally seek informed consent for participants for ethical reasons separately from the legal basis for data processing under GDPR, which is generally not informed consent. At a departmental level ethics committees generally do not have lawyers but they can advice researchers to obtain legal advice. Prior work has recommended researchers seek REB approval for legally risky research (such as Terms of Service violation) [42]. While REBs do not enforce the law, one key motivation for UK researchers to seek ethics committee review (seen in our interviews) is that they believe that this means that the institution has therefore approved the research project and so has to back them up if something then goes wrong [42].

H Additional Plots

See Figure 6, Figure 7, and Figure 8.

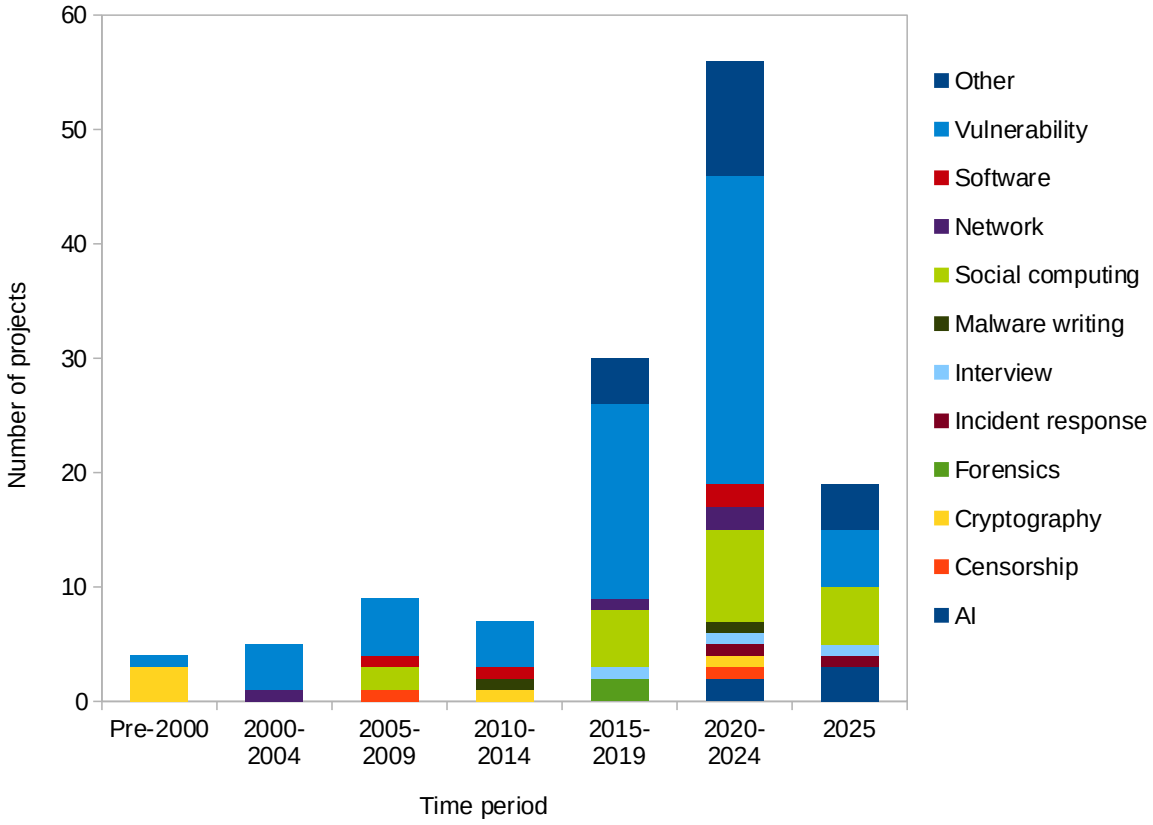

Figure 6: Number of projects of each research type per 5 year period. (Larger version of left plot of Figure 1.)

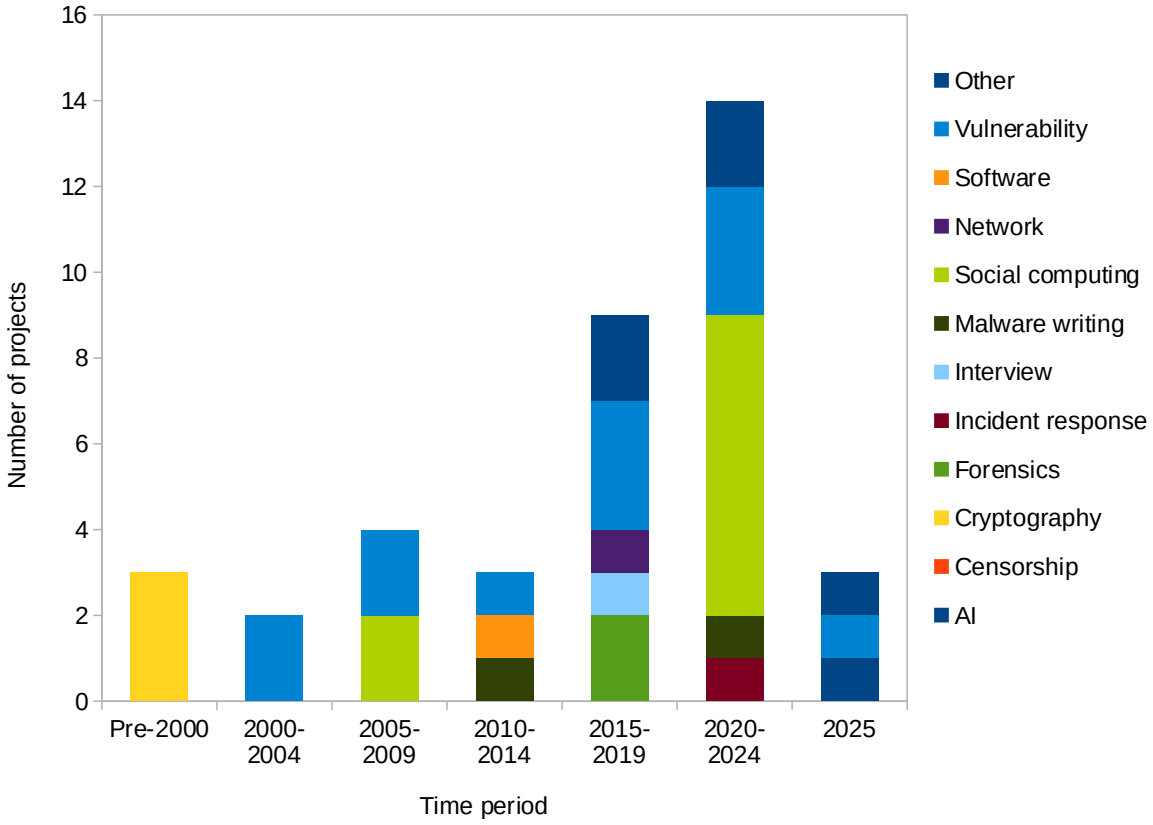

Figure 7: Number of projects involving a legal threat of each research type per 5 year period. (Larger version of right plot of Figure 1.)

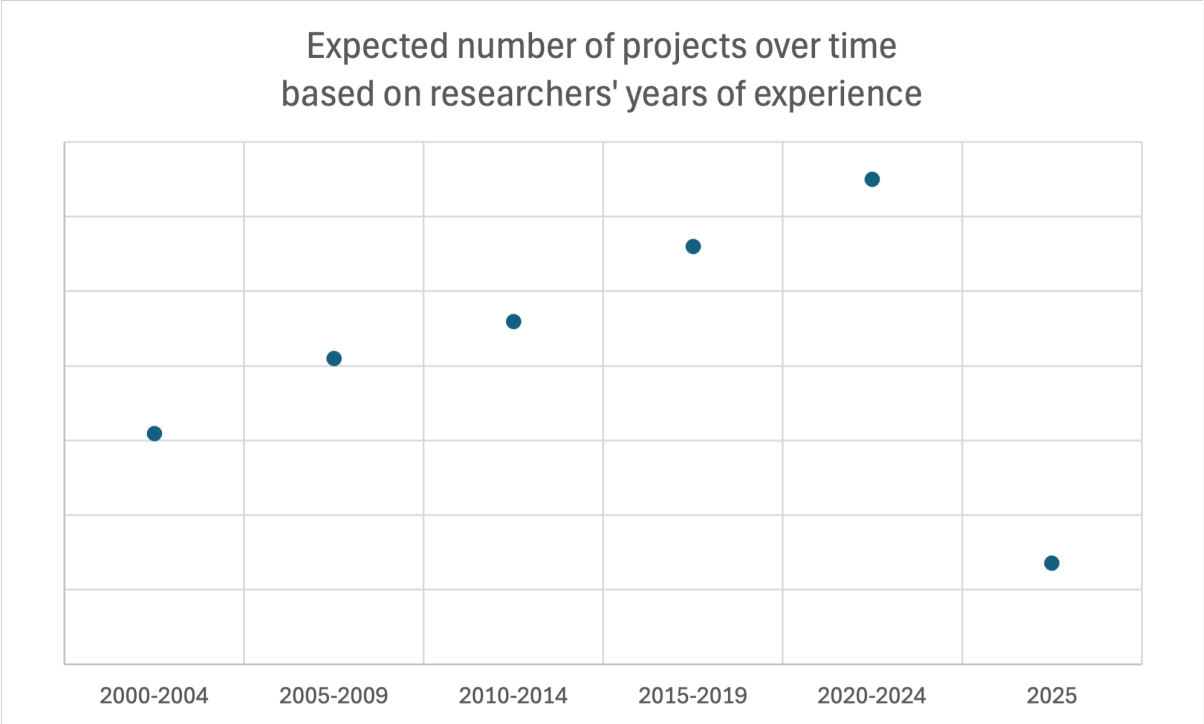

Figure 8: Expected number of projects over time based on researchers' years of experience, supposing that all researchers produce projects at an equal rate (intended as a rough reference for comparison)

References

[1] Judith Townend. Online chilling effects in England and Wales. *Internet Policy Review*, 3(2), April 3, 2014. DOI: 10.14763/2014.2.252.

[2] Jon Penney. Understanding chilling effects. *Minnesota Law Review 1451*, 106:80, 2022. DOI: 10.2139/ssrn.3855619.

[3] Jonathon W. Penney. *Chilling Effects*. Cambridge University Press, Cambridge, United Kingdom New York, NY, 2025. 286 pages. ISBN: 978-1-108-72530-9.

[4] Alexander Gamero-Garrido, Stefan Savage, Kirill Levchenko, and Alex C. Snoeren. Quantifying the pressure of legal risks on third-party vulnerability research. In *ACM SIGSAC Conference on Computer and Communications Security*, CCS '17, pages 1501–1513, October 2017. DOI: 10.1145/3133956.3134047.

[5] Qiu-Hong Wang, Ruibin Geng, and Seung Hyun Kim. Chilling effect of the enforcement of Computer Misuse Act: Evidence from publicly accessible hack forums. *Information Systems Research*, 35(3), September 19, 2023. DOI: 10.1287/isre.2019.0346.

[6] Sunoo Park and Kendra Albert. A researcher's guide to some legal risks of security research. Version 2, Harvard Cyberlaw Clinic, NYU Technology Law and Policy Clinic, Electronic Frontier Foundation (EFF), August 2, 2024, page 38. URL: https://clinic.cyber.harvard.edu/wp-content/uploads/2024/08/Security-Researchers-Guide-8-2-24.pdf. Archive-URL: https://perma.cc/43TZ-DW6G.

[7] Christopher Soghoian. Legal risks for phishing researchers. *eCrime Researchers Summit, eCrime 2008*, 2008. DOI: 10.1109/ECRIME.2008.4696971.

[8] Audrey Guinchard. The Computer Misuse Act 1990 to support vulnerability research? Proposal for a defence for hacking as a strategy in the fight against cybercrime. *Journal of Information Rights, Policy and Practice*, 2(2), March 27, 2017. DOI: 10.2139/ssrn.2946763.

[9] Daniel Etcovitch and Thyla van der Merwe. Coming in from the cold: A safe harbor from the CFAA and the DMCA §1201 for security researchers, Rochester, NY, June 1, 2018. DOI: 10.2139/ssrn.3055814.

[10] Andrew Sellars. Twenty years of web scraping and the Computer Fraud and Abuse Act. *Boston University Journal of Science and Technology Law*, 24, 2018. URL: https://scholarship.law.bu.edu/faculty_scholarship/465/ (visited on 02/05/2026). Archive-URL: https://perma.cc/7XTB-U2A2.

[11] Ram Shankar Siva Kumar, Jonathon Penney, Bruce Schneier, and Kendra Albert. Legal risks of adversarial machine learning research, ICML 2020 Workshop on Law & Machine Learning, June 29, 2020. DOI: 10.48550/arXiv.2006.16179.

[12] Simon McKay. Reforming the Computer Misuse Act 1990, Criminal Law Reform Now Network, 2020. URL: http://www.clrnn.co.uk/media/1018/clrnn-cma-report.pdf (visited on 04/27/2026). Archive-URL: https://perma.cc/JE9M-BTCG.

[13] Audrey Guinchard. The criminalisation of tools under the Computer Misuse Act 1990. The need to rethink cybercrime offences to effectively protect legitimate activities and deter cybercriminals. In *Rethinking Cybercrime: Critical Debates*, pages 41–61. Cham, 2021. DOI: 10.1007/978-3-030-55841-3_3.

[14] Madelyne Xiao, Andrew Sellars, and Sarah Scheffler. When anti-fraud laws become a barrier to computer science research. In *2025 Symposium on Computer Science and Law*, CSLAW '25. ACM, March 25, 2025. DOI: 10.1145/3709025.3712206.

[15] Jennifer Stisa Granick. Towards learning from losing Aaron Swartz. Stanford CIS. January 15, 2013. URL: https://cyberlaw.stanford.edu/blog/2013/01/towards-learning-losing-aaron-swartz/. Archive-URL: https://perma.cc/6PNS-4VS3.

[16] Aviel David Rubin. *Brave new ballot: The Battle to safeguard democracy in the age of electronic voting*. Crown, September 5, 2006. ISBN: 978-0-7679-2400-9.

[17] Alex Scroxton. Campaigners call for evidence to reform UK cyber laws. ComputerWeekly.com. July 31, 2024. URL: https://www.computerweekly.com/news/366599596/Campaigners-call-for-evidence-to-reform-UKs-cyber-laws. Archive-URL: https://perma.cc/3PJM-3MHA.

[18] Jennifer Mason. *Qualitative Researching*. SAGE, 2nd edition, 2002. 234 pages. ISBN: 0 7619 7427 X.

[19] Johnny Saldaña. *The coding manual for qualitative researchers*. 2013. ISBN: 978-1-4462-4736-5.

[20] Sunoo Park and Daniel R. Thomas. Supporting artefacts for: "Sticking their heads out above the parapets": Lived Experiences of Legal Risks in Research, June 2026. DOI: https://doi.org/10.15129/a0eae2f6-aa2a-4016-ab0a-3ed2ce7908d8.

[21] Eitan Frachtenberg and Rhody D. Kaner. Underrepresentation of women in computer systems research. *PLoS ONE*, 17(4):e0266439, April 6, 2022. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0266439. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC8985950/ (visited on 05/11/2026). Archive-URL: https://perma.cc/C38L-H8FQ.

[22] Overview: The state of the cybersecurity gender gap. New America. May 13, 2026. URL: https://www.newamerica.org/insights/new-ways-bring-women-and-through-cybersecurity-careers/overview-the-state-of-the-cybersecurity-gender-gap/ (visited on 05/11/2026). Archive-URL: https://perma.cc/8FGS-CEXZ.

[23] Virginia Braun and Victoria Clarke. One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative Research in Psychology*, July 3, 2021. DOI: 10.1080/14780887.2020.1769238.

[24] NVivo by Lumivero | Qualitative data analysis (QDA) software. Lumivero. URL: https://lumivero.com/products/nvivo/ (visited on 04/28/2026). Archive-URL: https://perma.cc/6NEU-HJRU.

[25] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proc. ACM Hum.-Comput. Interact.*, 3, CSCW, 2019. DOI: 10.1145/3359174.

[26] Andrew Gary Darwin Holmes. Researcher Positionality - A Consideration of Its Influence and Place in Qualitative Research - A New Researcher Guide. *Shanlax International Journal of Education*, 8(4):1–10, September 1, 2020. ISSN: 2582-1334. DOI: 10.34293/education.v8i4.3232.

[27] BOD 20-01: Develop and publish a Vulnerability Disclosure Policy. Cybersecurity & Infrastructure Security Agency (CISA). September 2, 2020. URL: https://www.cisa.gov/news-events/directives/bod-20-01-develop-and-publish-vulnerability-disclosure-policy (visited on 02/05/2026). Archive-URL: https://perma.cc/2DZ3-W7JF.

[28] Long comment regarding a proposed exemption under 17 U.S.C. 1201. In collaboration with Alex Rice. URL: https://www.copyright.gov/1201/2021/comments/Class%2013_InitialComments_HackerOne.pdf. Archive-URL: https://perma.cc/XPR3-CD9P.

[29] Office of Public Affairs. Department of Justice announces new policy for charging cases under the Computer Fraud and Abuse Act. May 19, 2022. URL: https://www.justice.gov/archives/opa/pr/department-justice-announces-new-policy-charging-cases-under-computer-fraud-and-abuse-act (visited on 02/05/2026). Archive-URL: https://perma.cc/3DDC-TPRU.

[30] 9-48.000 - Computer Fraud and Abuse Act, Department of Justice, May 19, 2022. URL: https://www.justice.gov/archives/opa/press-release/file/1507126/dl?inline. Archive-URL: https://perma.cc/XJJ6-DMEB.

[31] contingency fee. LII / Legal Information Institute. URL: https://www.law.cornell.edu/wex/contingency_fee (visited on 05/11/2026). Archive-URL: https://perma.cc/T39B-E2G2.

[32] Security Research Legal Defense Fund. URL: https://www.securityresearchlegaldefensefund.org/ (visited on 01/06/2026).

[33] Lorraine Kenny. Knight Institute to provide new legal support to researchers studying online platforms. Knight First Amendment Institute. November 13, 2023. URL: http://knightcolumbia.org/content/knight-institute-to-provide-new-legal-support-to-researchers-studying-online-platforms (visited on 04/16/2026). Archive-URL: https://perma.cc/7W6S-YYPX.

[34] Paul H. Rubin. More money into bad suits. *New York Times*, November 16, 2010. URL: https://www.nytimes.com/roomfordebate/2010/11/15/investing-in-someone-elses-lawsuit/more-money-into-bad-suits (visited on 02/05/2026). Archive-URL: https://perma.cc/S9S3-5JCM.

[35] Marleen Weulen Kranenbarg, Thomas J. Holt, and Jeroen van der Ham. Don't shoot the messenger! A criminological and computer science perspective on coordinated vulnerability disclosure. *Crime Science*, 7(1):16, November 19, 2018. ISSN: 2193-7680. DOI: 10.1186/s40163-018-0090-8.

[36] Kevin Beaumont. Microsoft's stance on zero day exploits is a dumpster fire of their own making. May 28, 2026. URL: https://doublepulsar.com/microsofts-stance-on-zero-day-exploits-is-a-dumpster-fire-of-their-own-making-0946117940a4 (visited on 06/01/2026). Archive-URL: https://perma.cc/BAX3-GMFF.

[37] Kendra Albert. Everything old Is new again: Legal restrictions on vulnerability disclosure on bug bounty platforms. *Usenix Security Symposium 2025*, 2025. URL: https://www.usenix.org/conference/usenixsecurity25/presentation/albert.

[38] Shayne Longpre et al. Position: In-house evaluation is not enough. Towards robust third-party evaluation and flaw disclosure for general-purpose AI. In *ICML*, July 15, 2025. URL: https://icml.cc/virtual/2025/poster/40170 (visited on 04/27/2026). Archive-URL: https://perma.cc/773K-EWT2.

[39] Alex Abdo et al. Better access: Data for the common good. Knight-Georgetown Institute. November 6, 2025. URL: https://kgi.georgetown.edu/research-and-commentary/better-access/. Archive-URL: https://perma.cc/GXJ2-6R8M.

[40] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. ZMap: Fast Internet-wide scanning and its security applications. In *Proceedings of the 22nd USENIX Security Symposium*, 2013. ISBN: 978-1-931971-03-4.

[41] David Dittrich and Erin Kenneally. The Menlo report: Ethical principles guiding information and communication technology research, U.S. Department for Homeland Security, August 2012, pages 1–13. URL: https://www.caida.org/publications/papers/2012/menlo_report_actual_formatted/. Archive-URL: https://perma.cc/7R2N-PW96.

[42] Daniel R. Thomas, Sergio Pastrana, Alice Hutchings, Richard Clayton, and Alastair R. Beresford. Ethical issues in research using datasets of illicit origin. In *Internet Measurement Conference*, IMC, London, UK. ACM, November 2017. DOI: 10.1145/3131365.3131389.

[43] Harshini Sri Ramulu, Helen Schmitt, Bogdan Rerich, Rachel Gonzalez Rodriguez, Tadayoshi Kohno, and Yasemin Acar. [Extended] Ethics in Computer Security Research: A Data-Driven Assessment of the Past, the Present, and the Possible Future. In *ACM Conference on Computer and Communications Security (CCS '25)*, September 2025. DOI: 10.1145/3719027.3765071. URL: http://arxiv.org/abs/2509.09351.

[44] Jessica Vitak, Nicholas Proferes, Katie Shilton, and Zahra Ashktorab. Ethics Regulation in Social Computing Research: Examining the Role of Institutional Review Boards. *Journal of Empirical Research on Human Research Ethics*, 12(5):372–382, December 2017. DOI: 10.1177/1556264617725200.

[45] Simson L. Garfinkel and Lorrie Faith Cranor. Institutional Review Boards and your research. *Communications of the ACM*, June 1, 2010. DOI: 10.1145/1743546.1743563.

[46] Office for Human Research Protections. Institutional Review Board written procedures: Guidance for Institutions and IRBs. 2025. URL: https://www.hhs.gov/ohrp/regulations-and-policy/guidance/institutional-issues/institutional-review-board-written-procedures/index.html (visited on 05/11/2026). Archive-URL: https://perma.cc/9XXV-NWEH.

[47] National Institutes of Health. Certificates of Confidentiality (CoC). August 10, 2024. URL: https://grants.nih.gov/policy-and-compliance/policy-topics/human-subjects/coc (visited on 03/09/2026). Archive-URL: https://perma.cc/USS7-CT99.

[48] National Institutes of Health. CoCs for research not funded by NIH. August 22, 2025. URL: https://grants.nih.gov/policy-and-compliance/policy-topics/human-subjects/coc/for-research-not-funded-by-nih (visited on 03/09/2026). Archive-URL: https://perma.cc/SCD3-HBTZ.

[49] Kendra Albert and James Grimmelmann. Do the right thing. *Communications of the ACM*, 66(5):18–20, May 2023. DOI: 10.1145/3587825.

[50] UK Research {and} Innovation. Trusted research and innovation. October 13, 2025. URL: https://www.ukri.org/manage-your-award/good-research-resource-hub/trusted-research-and-innovation/ (visited on 04/04/2025). Archive-URL: https://perma.cc/2HND-B9PY.

[51] National Protective Security Authority (NPSA). Trusted Research guidance for academia. July 2, 2024. URL: https://www.npsa.gov.uk/trusted-research-academia (visited on 04/04/2025). Archive-URL: https://web.archive.org/web/20250517085013/https://www.npsa.gov.uk/trusted-research-academia.

[52] Commonwealth {and} Development Office Foreign. Academic Technology Approval Scheme (ATAS). January 8, 2025. URL: https://www.gov.uk/guidance/academic-technology-approval-scheme (visited on 04/04/2025). Archive-URL: https://perma.cc/PY5J-4QHL.

[53] Participation in EU calls with ownership and control restrictions. Funding guidance Version 2.0, European Commission, January 1, 2026. URL: https://ec.europa.eu/info/funding-tenders/opportunities/docs/2021-2027/common/guidance/guidance-on-participation-in-eu-calls-with-ownership-and-control-restrictions_en.pdf (visited on 05/11/2026). Archive-URL: https://perma.cc/3YVB-YTWH.

[54] Kayleigh-Ann Clegg. Research Implications of Increased Political Oversight in the United States. *Journal of Medical Internet Research*, 27:e83499, September 19, 2025. ISSN: 1439-4456. DOI: 10.2196/83499. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC12451217/ (visited on 05/11/2026).

[55] Lisa Larrimore Ouellette. Government research funding and the First Amendment. January 27, 2026. URL: https://law.yale.edu/yls-today/yale-law-school-events/government-research-funding-and-first-amendment-lisa-larrimore (visited on 05/11/2026). Archive-URL: https://perma.cc/SW4X-NV4H.

[56] Dan Garisto. Trump order gives political appointees vast powers over research grants. *Nature*, 644(8077):585–586, August 8, 2025. ISSN: 1476-4687. DOI: 10.1038/d41586-025-02557-z. URL: https://www.nature.com/articles/d41586-025-02557-z (visited on 05/11/2026). Archive-URL: https://perma.cc/JDC2-396X.

[57] Timothy Snyder. *On Tyranny: Twenty Lessons from the Twentieth Century*. Crown, February 28, 2017. 128 pages. ISBN: 978-0-8041-9011-4.

[58] Jared L. Hubbard. Brief of Voatz, Inc., as Amicus Curiae in support of respondent in affirmance of the decision below, Supreme Court of the United States, September 3, 2020. URL: https://www.supremecourt.gov/DocketPDF/19/19-783/153062/20200903122434600_Voatz%20Amicus%20Brief.pdf. Archive-URL: https://perma.cc/U8MR-2DRM.

[59] MSRC Team. A shared responsibility: Protecting customers through Coordinated Vulnerability Disclosure. May 27, 2026. URL: https://www.microsoft.com/en-us/msrc/blog/2026/05/a-shared-responsibility-protecting-customers-through-coordinated-vulnerability-disclosure (visited on 06/01/2026). Archive-URL: https://perma.cc/QNG7-DTLR.

[60] Jim Baker. Rethinking encryption. *Lawfare*, January 18, 2023. URL: https://www.lawfaremedia.org/article/rethinking-encryption. Archive-URL: https://perma.cc/FT6U-HXHU.

[61] Biden Administration. National Cybersecurity Strategy, Whitehouse, March 2023. URL: https://bidenwhitehouse.archives.gov/wp-content/uploads/2023/03/National-Cybersecurity-Strategy-2023.pdf (visited on 05/11/2026). Archive-URL: https://perma.cc/GV9S-4BKN.

[62] It's time to act: Open your eyes to the imminent risk to your economic security: NCSC Annual Review 2025, National Cyber Security Centre, 2025. URL: https://www.ncsc.gov.uk/sites/default/files/documents/ncsc-annual-review-2025.pdf. Archive-URL: https://perma.cc/72NY-PPPJ.

[63] NCSC Annual Review 2024, National Cyber Security Centre (NCSC), 2024. URL: https://www.ncsc.gov.uk/sites/default/files/documents/NCSC_Annual_Review_2024.pdf. Archive-URL: https://perma.cc/8J3M-FC48.

[64] NSA Media Relations. NSA joins CISA and Others in Releasing the Cybersecurity Information Sheet "Establishing a. National Security Agency/Central Security Service. July 15, 2026. URL: http://www.nsa.gov/Press-Room/Press-Releases-Statements/Press-Release-View/Article/4546549/nsa-joins-cisa-and-others-in-releasing-the-cybersecurity-information-sheet-esta/ (visited on 09/10/2026).

[65] VDP Platform Team. Vulnerability Disclosure Policy (VDP) Platform | CISA. URL: https://www.cisa.gov/resources-tools/services/vulnerability-disclosure-policy-vdp-platform (visited on 09/10/2026).

[66] Information Technology Laboratory Computer Security Division. Vulnerability Disclosure Guidelines | CSRC | CSRC. CSRC | NIST. February 4, 2021. URL: https://csrc.nist.gov/projects/vulnerability-disclosure-guidelines (visited on 09/10/2026).

[67] Riana Pfefferkorn. America's anti-hacking laws pose a risk to national security. Brookings. September 7, 2021. URL: https://www.brookings.edu/techstream/americas-anti-hacking-laws-pose-a-risk-to-national-security/. Archive-URL: https://perma.cc/LYV2-224T.

[68] United States Government. Vulnerabilities Equities Policy and Process for the United States Government. November 15, 2017. URL: https://www.whitehouse.gov/sites/whitehouse.gov/files/images/External%20-%20Unclassified%20VEP%20Charter%20FINAL.PDF (visited on 06/28/2018).

[69] Big Sleep team. From Naptime to Big Sleep: Using large language models to catch vulnerabilities in real-world code. November 1, 2024. URL: https://projectzero.google/2024/10/from-naptime-to-big-sleep.html (visited on 05/27/2026). Archive-URL: https://perma.cc/322C-RR4R.

[70] Anthropic's Frontier Red Team. red.anthropic.com. May 2026. URL: https://red.anthropic.com/ (visited on 05/27/2026). Archive-URL: https://perma.cc/VFJ8-LCZ3.

[71] Kyle Orland. Amid Mythos' hyped cybersecurity prowess, researchers find GPT-5.5 is just as good. Ars Technica. May 1, 2026. URL: https://arstechnica.com/ai/2026/05/amid-mythos-hyped-cybersecurity-prowess-researchers-find-gpt-5-5-is-just-as-good/. Archive-URL: https://web.archive.org/web/20260504193441/https://arstechnica.com/ai/2026/05/amid-mythos-hyped-cybersecurity-prowess-researchers-find-gpt-5-5-is-just-as-good/.

[72] Rock Lambros. AI vulnerability discovery: Mythos Is the headline. Not the story. April 14, 2026. URL: https://www.rockcybermusings.com/p/ai-vulnerability-discovery-mythos. Archive-URL: https://perma.cc/D3RK-8D64.

[73] Dan Goodin. Mozilla says 271 vulnerabilities found by Mythos have "almost no false positives". Ars Technica. May 7, 2026. URL: https://arstechnica.com/information-technology/2026/05/mozilla-says-271-vulnerabilities-found-by-mythos-have-almost-no-false-positives/. Archive-URL: https://web.archive.org/web/20260515040049/https://arstechnica.com/information-technology/2026/05/mozilla-says-271-vulnerabilities-found-by-mythos-have-almost-no-false-positives/.

[74] Google Threat Intelligence Group. Adversaries leverage AI for vulnerability exploitation, augmented operations, and initial access. Google Cloud Blog. May 11, 2026. URL: https://cloud.google.com/blog/topics/threat-intelligence/ai-vulnerability-exploitation-initial-access. Archive-URL: https://perma.cc/64N5-HN3W.

[75] Sead Fadilpašić. Anthropic claims Mythos has found over ten thousand major security vulnerabilities. TechRadar. May 25, 2026. URL: https://www.techradar.com/pro/security/after-one-month-most-partners-have-each-found-hundreds-of-critical-or-high-severity-vulnerabilities-anthropic-claims-mythos-has-found-over-ten-thousand-major-security-vulnerabilities-across-the-most-systemically-important-software-in-the-world. Archive-URL: https://perma.cc/HPK8-WLFH.

[76] Nicholas Carlini et al. Assessing Claude Mythos Preview's cybersecurity capabilities. April 7, 2026. URL: https://red.anthropic.com/2026/mythos-preview/ (visited on 05/27/2026).

[77] Anthropic Frontier Red Team. Anthropic's coordinated vulnerability disclosure dashboard. May 22, 2026. URL: https://red.anthropic.com/2026/cvd/.

[78] NACDL - CFAA Background. National Association of Criminal Defense Lawyers. July 14, 2022. URL: https://www.nacdl.org/Content/CFAABackground. Archive-URL: https://perma.cc/E7S9-N4N9.

[79] Stefan Fafinski. Computer Misuse: The Implications of the Police and Justice Act 2006. *The Journal of Criminal Law*, 72(1):53–66, February 1, 2008. ISSN: 0022-0183. DOI: 10.1350/jcla.2008.72.1.477.

[80] Kat Sommer and Verona Johnstone-Hulse. Now is the time to CyberUp – making the Computer Misuse Act fit for the 21st century. March 28, 2025. URL: https://www.nccgroup.com/now-is-the-time-to-cyberup-making-the-computer-misuse-act-fit-for-the-21st-century/. Archive-URL: https://perma.cc/B5WR-CUQV.

[81] Cindy Cohn. Aaron's Law reintroduced: CFAA didn't fix itself. Electronic Frontier Foundation. April 29, 2015. URL: https://www.eff.org/deeplinks/2015/04/aarons-law-reintroduced-cfaa-didnt-fix-itself (visited on 05/11/2026). Archive-URL: https://perma.cc/99CK-EBKM.

[82] Alex Scroxton. UK government pledges to rewrite Computer Misuse Act. Computer Weekly. December 3, 2025. URL: https://www.computerweekly.com/news/366635624/UK-government-pledges-to-rewrite-Computer-Misuse-Act. Archive-URL: https://perma.cc/8EBJ-Y9VG.

[83] Nat Meysenburg. Cybersecurity research should not be a crime. New America. May 13, 2026. URL: https://www.newamerica.org/insights/cybersecurity-research-should-not-be-a-crime/. Archive-URL: https://perma.cc/DTY5-8A4Y.

[84] TLPC Advocates for expansion of security research exemption to Section 1201 of the DMCA – Samuelson-Glushko Technology Law & Policy Clinic. April 19, 2021. URL: https://tlpc.colorado.edu/tlpc-advocates-for-expansion-of-security-research-exemption-to-section-1201-of-the-dmca/. Archive-URL: https://perma.cc/K88T-ZMGE.

[85] Stan Adams. Once more, with feeling: Security research should not be chilled by uncertainty in copyright law. Center for Democracy and Technology. December 15, 2020. URL: https://cdt.org/insights/once-more-with-feeling-security-research-should-not-be-chilled-by-uncertainty-in-copyright-law/. Archive-URL: https://perma.cc/2PFR-3L68.

[86] Andrew Crocker and Naomi Gilens. Van Buren - EFF security researchers amicus brief. Electronic Frontier Foundation. July 8, 2020. URL: https://www.eff.org/document/van-buren-eff-security-researchers-amicus-brief. Archive-URL: https://perma.cc/5DJW-4XNJ.

[87] J. Alex Halderman, Blake E. Reid, Paul K. Ohm, Harry A. Surden, and Bernthal J. Brad. Before the Copyright Office of the Library of Congress: In the matter of exemption to prohibition on circumvention of copyright protection systems for access control technologies. Docket No. RM 2008 8. December 2, 2008. URL: https://tlpc.colorado.edu/wp-content/uploads/2021/04/halderman-reid.pdf. Archive-URL: https://perma.cc/FG27-C7V2.

[88] Harley Geiger. NIST Cyber Framework To Include Coordinated Vuln Disclosure Processes | Rapid7 Blog. Rapid7. December 19, 2017. URL: https://www.rapid7.com/blog/post/2017/12/19/nist-cyber-framework-revised-to-include-coordinated-vuln-disclosure-processes/ (visited on 09/10/2026).

[89] CyberUp Campaign. April 16, 2026. URL: https://www.cyberupcampaign.com (visited on 05/19/2026). Archive-URL: https://perma.cc/WW7T-W385.

[90] Joseph Lorenzo Hall and Stan Adams. Taking the pulse of hacking: A risk basis for security reearch, Center for Democracy & Technology (CDT), March 27, 2018. URL: https://cdt.org/wp-content/uploads/2018/04/2018-03-27-Risk-Basis-for-Security-Research-FNL.pdf. Archive-URL: https://perma.cc/4CHM-MKMK.

[91] Omer Akgul, Taha Eghtesad, Amit Elazari, Omprakash Gnawali, Jens Grossklags, Michelle L. Mazurek, Daniel Votipka, and Aron Laszka. Bug hunters' perspectives on the challenges and benefits of the bug bounty ecosystem. *Usenix Security Symposium 2023*, 2023. URL:

https://www.usenix.org/system/files/usenixsecurity23-akgul.pdf.

[92] Temima Hrle, Mary Milad, Jingjie Li, and Daniel Woods. "Just a tool, until you stab someone with it": Exploring Reddit users' questions and advice on the legality of port scans. In *Proceedings of the 2024 European Symposium on Usable Security*, EuroUSEC '24, pages 322–336, New York, NY, USA. Association for Computing Machinery, November 20, 2024. ISBN: 979-8-4007-1796-3. DOI: 10.1145/3688459.3688469. (Visited on 04/27/2026).

[93] Aaron J. Burstein. Conducting cybersecurity research legally and ethically. In *USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET)*, 2008. URL: https://www.usenix.org/event/leet08/tech/full_papers/burstein/burstein_html. Archive-URL: https://perma.cc/Q9BM-MDV5.

[94] Ivan Evtimov, David O'Hair, Earlence Fernandes, Ryan Calo, and Tadayoshi Kohno. Is tricking a robot hacking? *Berkeley Technology Law Journal*, 34(3), 2020. ISSN: 1556-5068. DOI: 10.15779/Z38M32N99Q.

[95] Paul Ohm, Douglas C. Sicker, and Dirk Grunwald. Legal issues surrounding monitoring during network research. *7th ACM SIGCOMM conference on Internet measurement*:141–148, 2007. DOI: 10.1145/1298306.1298307.

[96] Riana Pfefferkorn. Shooting the messenger: Remediation of disclosed vulnerabilities as CFAA "Loss". *Richmond Journal of Law & Technology*, 29(1):89–164, November 30, 2022. ISSN: 1091-7322. DOI: 10.2139/ssrn.4224982.

[97] Aaron J. Burstein. Amending the ECPA to enable a culture of cybersecurity research. *Harvard Journal of Law & Technology*, 22(1):167–222, 2008.

[98] Periwinkle Doerfler, Andrea Forte, Emiliano De Cristofaro, Gianluca Stringhini, Jeremy Blackburn, and Damon McCoy. "I'm a Professor, which isn't usually a dangerous job": Internet-facilitated harassment and its impact on researchers. *Proceedings of the ACM on Human-Computer Interaction*, 5:341:1–341:32, CSCW2, October 18, 2021. DOI: 10.1145/3476082.

[99] Alice Marwick et al. Risky research: An AoIR guide to researcher protection and safety, The Association of Internet Researchers, 2025. URL: https://aoir.org/riskyresearchguide/. Archive-URL: https://perma.cc/6KVF-Y3MY.

[100] Sarah Gilbert, J. Nathan Matias, and Floor Fiers. Sustainably managing risks and threats to independent research, Open Science Framework, March 27, 2024. URL: https://osf.io/ex4p8 (visited on 04/04/2024).

[101] Harshini Sri Ramulu, Anna Lena Rotthaler, Jost Rossel, Rachel Gonzalez Rodriguez, Dominik Wermke, Sascha Fahl, Tadayoshi Kohno, Juraj Somorovsky, and Yasemin Acar. Poster: Computer security researchers' experiences with vulnerability disclosures. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security*, CCS '25, pages 4779–4781, New York, NY, USA. Association for Computing Machinery, November 22, 2025. ISBN: 979-8-4007-1525-9. DOI: 10.1145/3719027.3760723. (Visited on 02/04/2026).

[102] Giovane C. M. Moura and John Heidemann. Vulnerability disclosure considered stressful. *SIGCOMM Comput. Commun. Rev.*, 53(2):2–10, July 19, 2023. ISSN: 0146-4833. DOI: 10.1145/3610381.3610383.

[103] David Basin and Srdjan Capkun. The research value of publishing attacks. *Commun. ACM*, 55(11):22–24, November 1, 2012. ISSN: 0001-0782. DOI: 10.1145/2366316.2366324.

[104] Tadayoshi Kohno, Yasemin Acar, and Wulf Loh. Ethical frameworks and computer security trolley problems: Foundations for conversations. 2023. URL: https://securityethics.cs.washington.edu/ (visited on 01/06/2026). Full version of USENIX Security 2023 paper.

[105] Amit Elazari. Hacking the law: Are bug bounties a true safe harbor? In Enigma 2018. USENIX, 2018. URL: https://www.usenix.org/conference/enigma2018/presentation/elazari.

[106] Amit Elazari. Private ordering shaping cybersecurity policy: The case of bug bounties. In *Rewired: Cybersecurity Governance*. Wiley, Rochester, NY, April 2019. ISBN: 978-1-118-88821-6. URL: https://papers.ssrn.com/abstract=3161758 (visited on 05/26/2026).

[107] Sunoo Park and Daniel R. Thomas. legal-risk-resources. June 3, 2026. URL: https://codeberg.org/nooske/legal-risk-resources.

[108] Electronic Frontier Foundation. URL: https://www.eff.org/ (visited on 01/06/2026). Archive-URL: https://perma.cc/43W3-WUC9.

[109] Knight First Amendment Institute. Knight First Amendment Institute. URL: http://knightcolumbia.org/ (visited on 01/06/2026). Archive-URL: https://perma.cc/V35B-8L4F.

[110] Kendra Albert and Andrew Sellars. Albert Sellars. Albert Sellars. URL: https://www.albertsellars.law (visited on 01/06/2026). Archive-URL: https://perma.cc/VCL4-WHBX.

[111] Richard Hanstock. pwn.legal: Community edition. URL: https://help.pwn.legal/ (visited on 05/21/2026). Archive-URL: https://perma.cc/43T9-3XVZ.

[112] BU/MIT Student Innovations Law Clinic. URL: https://www.bu.edu/law/experiential-learning/clinics/bu-mit-student-innovations-law-clinic/ (visited on 01/06/2026). Archive-URL: https://perma.cc/L8AA-PFZL.

[113] Samuelson-Glushko Technology Law & Policy Clinic | Colorado Law | University of Colorado Boulder. URL: https://www.colorado.edu/law/academics/clinics/samuelson-glushko-technology-law-policy-clinic (visited on 01/06/2026). Archive-URL: https://perma.cc/W8P5-9YND.

[114] Cyberlaw Clinic. Harvard Law School. URL: https://hls.harvard.edu/clinics/in-house-clinics/cyberlaw-clinic/ (visited on 01/06/2026). Archive-URL: https://perma.cc/L5P4-BBEE.

[115] NYU Technology Law & Policy Clinic. NYU School of Law. URL: https://www.law.nyu.edu/academics/clinics/tech-law-policy (visited on 05/11/2026). Archive-URL: https://perma.cc/9YCG-Y6HU.

[116] disclose.io. The disclose.io Project. URL: https://disclose.io/ (visited on 02/08/2024). Archive-URL: https://perma.cc/4F6K-24QU.

[117] Front Page - I Am The Cavalry. URL: https://iamthecavalry.org/ (visited on 05/11/2026). Archive-URL: https://perma.cc/N78H-GBNR.

[118] AHA! Is Austin Hackers Anonymous. AHA! URL: http://takeonme.org/ (visited on 05/11/2026). Archive-URL: https://perma.cc/V7VA-7FT9.

[119] threats.disclose.io. Research threats: Legal threats against security researchers, January 25, 2025. URL: https://github.com/disclose/research-threats. Archive-URL: https://github.com/disclose/research-threats.

[120] Louis Barclay. How to deal with receiving a cease-and-desist letter from Big Tech. 12 Challenges. January 30, 2024. URL: https://12challenges.substack.com/p/how-to-deal-with-receiving-a-cease (visited on 01/30/2024). Archive-URL: https://perma.cc/PC3Y-DVTC.

[121] Consortium. Tips for scholars engaged in public-facing communication about elections, Researcher Support Consortium, 2024. URL: https://researchersupport.org/wp-content/uploads/2024/07/Tips_for_Safer_Communication.pdf (visited on 08/06/2024). Archive-URL: https://perma.cc/Z6NP-ZG9K.

[122] Open Rights Group (ORG). URL: https://www.openrightsgroup.org/ (visited on 06/01/2026). Archive-URL: https://perma.cc/GTC5-J84B.

[123] Coalition for Independent Technology Research (CITR). URL: https://independenttechresearch.org/ (visited on 06/01/2026). Archive-URL: https://perma.cc/2BPN-DDAM.

[124] Content {and} Technology European Commission Directorate-General for Communications Networks. Delegated act on data access under the Digital Services Act (DSA). July 2, 2025. URL: https://digital-strategy.ec.europa.eu/en/library/delegated-act-data-access-under-digital-services-act-dsa (visited on 05/19/2026). Archive-URL: https://perma.cc/MC4Q-VVGQ.

[125] Ben Adida et al. Statement on legal impediments to cybersecurity research, Filing to US Copyright Office about the security research exemptions., May 21, 2015. URL: https://cdn.loc.gov/copyright/1201/2015/hearing-exhibits/Cybersec-statement-5-21-15.pdf. Archive-URL: https://perma.cc/3PXR-EZBC.

[126] Supreme Court of the United States 19-783 Van Buren v. United States, URL: https://www.supremecourt.gov/opinions/20pdf/19-783_k53l.pdf.

[127] UKRI policy and guidance on the governance of good research practice. October 17, 2025. URL: https://www.ukri.org/councils/esrc/guidance-for-applicants/research-ethics-guidance/our-policy-and-guidelines-for-good-research-conduct/ (visited on 02/25/2026). Archive-URL: https://perma.cc/QHM9-HYZL.

[128] The Concordat to Support Research Integrity – UKCORI. 2025. URL: https://ukcori.org/research-integrity-concordat/ (visited on 02/25/2026). Archive-URL: https://perma.cc/E6WG-KP99.

[129] Research organisations and research ethics committees. UK Research and Innovation (UKRI). February 13, 2026. URL: https://www.ukri.org/councils/esrc/guidance-for-applicants/research-ethics-guidance/research-organisations-and-research-ethics-committees-our-principles-research-ethics-committees/ (visited on 02/25/2026). Archive-URL: https://perma.cc/U7W5-LR6J.

[130] Research involving human participants policy – Grant funding. Wellcome. August 31, 2025. URL: https://wellcome.org/research-funding/guidance/policies-grant-conditions/research-involving-human-participants-policy (visited on 02/25/2026). Archive-URL: https://perma.cc/9GPU-TSJT.

[131] Stewart Carswell. Ethical review of research. May 24, 2019. URL: https://www.cst.cam.ac.uk/local/policy/ethics (visited on 02/25/2026). Archive-URL: https://perma.cc/U2ZH-AJX3.

[132] University of Cambridge policy on the ethical conduct of research involving human participants or personal data. URL: https://www.research-services.admin.cam.ac.uk/sites/default/files/research-ethics-policy.pdf (visited on 02/25/2026). Archive-URL: https://perma.cc/6CZG-XEB3.

[133] University Ethics Committee. Code of Practice on investigations involving human beings. Eighth Edition, University of Strathclyde, February 2017, pages 1–50. URL: https://www.strath.ac.uk/media/ps/rkes/Code_of_Practice_eighth_Feb17.pdf (visited on 02/25/2026). Archive-URL: https://perma.cc/7SCV-RZKV.