Why Ransomware Explains Itself: Usable Security and Operational Dependency

Mark Quinlan, and Aaron Ceross

University College London, University of Birmingham

Abstract. Adams and Sasse's *Users Are Not the Enemy* propelled the treatment of users into the foreground of security research and practice. Yet twenty-five years on, users locked out of platform accounts routinely encounter interactions far less navigable than those offered by ransomware operators seeking payment. This paper argues that this disparity is structural, not accidental. We introduce *procedural respect*, defined as whether an interaction is designed to succeed rather than merely to exist, independent of its moral legitimacy, and identify three conditions that render its absence institutionally rational: decoupled value, exit foreclosure, and opacity-as-protection. Our analysis yields a diagnostic question: does the operator need the interaction to succeed? Where the answer is yes, usable security follows; where it is no, procedural opacity should be expected. Sasse's philosophy remains the normative benchmark; our contribution is to specify where structural conditions permit its realisation, and what must change where they do not.

1 Introduction

Let us consider two moments of constrained access. In the first, you attempt to log into a familiar platform and are informed that *suspicious activity* has led to you being locked out. To restore access, you are invited to verify your identity, yet every submitted document is rejected without explanation. You find no clear escalation path and no indication of how long the process will last. The system informs you that something has gone wrong, but it does not tell you what that something is [14]. You wait, and the silence becomes a message in its own right [39].

In the second moment, ransomware seizes control of your files. A stark message fills the screen and announces what has occurred. It explains why your files are inaccessible [17,40], specifies what and how much the attackers demand [26], provides a deadline, and offers a communication channel. Instructions are explicit [18,23]. The path to resolution (however coercive) is clear. You may even

This work is licensed under a Creative Commons "Attribution-NonCommercial 4.0 International" license. To view a copy of this license visit https://creativecommons.org/licenses/by-nc/4.0/deed.en. ©2026 Copyright held by the owner/author(s).

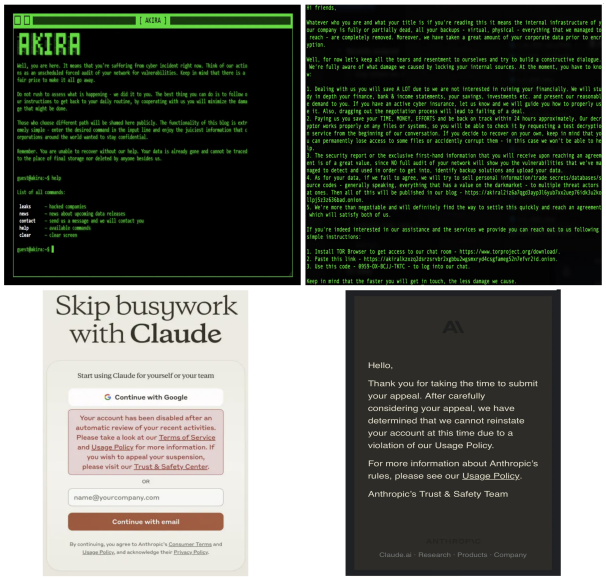

Fig. 1. Constrained-access communications. Top row: A ransomware portal. Bottom row: A Platform lockout and appeal rejection message. The ransomware operator explains, a follow-up screen would simply be the return to status quo ante; the platform offers little at the start of the interaction, even less when failure is communicated.

test whether decryption is possible before committing [8]. The situation is frightening, yet structurally navigable[1].

Are these scenarios equivalent then? Of course not, you would say. For one, the first is lawful service provision, the second criminal extortion. But this moral asymmetry is what makes the comparison interesting. For if the quality of such an interaction tracked any given form of common moral alignment, the platform provider would take usable security, at least through its interactions, much more seriously than the criminal. But it does not, and these moments suggest instead that interactional quality follows operational dependency, since clarity appears only where the restricting party needs the user to succeed.

Furthermore, both are examples of usable security, as they involve the same interactional security-based relationship: a powerful actor restricts access, dictates conditions, and controls the path toward resolution. The user's goal in both cases is identical, namely to return to their status quo *ante*. What differs is how each system treats the interaction through which resolution might occur. The ransomware operator invests in that interaction; the platform does not.

[1] The user has become so used to these digressions into their daily online activity that malicious actors have themselves taken advantage, by designing elaborate duplicates of Meta's "Suspension Pending Review" lock-out. https://malwaretips.com/blogs/meta-suspension-pending-review/.

Hence we arrive at a comparison that can show a criminal enterprise communicating more clearly than a system claiming to safeguard users. This outcome stands in sharp tension with Adams and Sasse's seminal Users Are Not the Enemy, which we use here as a foundational reference point for examining where and why its principles are realised (or systematically abandoned) in contemporary platform security practice. In their account, users who circumvent controls, forget passwords, or fall for social engineering are not deficient; rather, they are responding rationally to systems that impose unreasonable burdens [1]. Durable arguments that anchor a research tradition, whose principles have been verified [3, 30], taught [15, 32], applied as frameworks [24, 29, 31], provoked direct polemics [38], and have resulted in genuine progress in furthering the goals of usable security. At this point, the mantra that users are not the enemy can almost be seen as a guiding philosophy within the wider Human-Computer Interaction field, and yet, in practice, operations analogous to our example, reasonably frequent interactions where users are most dependent on institutional support, routinely exhibit the poorest design quality. The question is why.

Any worthy philosophy has provoked direct polemics, and UANTE[2] is no different. Vidyaraman et al., for instance, advance a contrary design stance, "the user *is* the enemy" which treats users as adversarial system components whose actions must be constrained and penalised rather than supported [38]. Their counter-position illustrates that the debate has long been framed in moral or attitudinal terms, that they represent a failure of knowledge transfer or institutional will. Our arguments in this paper depart from this framing by showing that design outcomes follow from *structural dependency*, not professed philosophy.

What we mean by this dependency is that systems implement the ideals of Usable Security when doing so serves operational goals; they do not when operational goals are served by neglect. In our example, the ransomware operator, who needs the victim to navigate a payment process successfully, inadvertently implements the philosophy (or at least one similar) more faithfully than the platform, which we argue faces little to no penalty for user failure during recovery. But neither actor is motivated by the philosophy itself. One is constrained by circumstance to behave as if users matter; the other is not.

Arguments such as these require careful delimitation. In our example, we examined only the interactional surface that emerges when access is restricted, the communications, pathways, and feedback mechanisms (in practice we reviewed studies of splash screens, payment portals, and operator behaviour [8,17,26,40]) through which a constrained user might seek resolution [3]. We did not compare the moral legitimacy, technical sophistication, nor societal consequences of ransomware and platform operations. Nor did or will we be treating the user as an enemy. Our method in this comparison are based on defamiliarisation [4,33], which is simply using an unexpected comparison to make visible what familiar-

[2] Our acronym for "Users Are Not the Enemy" in the remainder of this paper.

[3] We do not seek in this paper to examine technical/social outcomes, this is left to many other studies on this topic, e.g, Filiz et al [12].

ity has normalised. The ransomware case renders strange what platform security has made routine. Using this method allows us to synthesise existing studies to examine the interactional surface without presenting new empirical data.

To contextualize our work on the interactional surface we introduce the concept of *procedural respect*. Procedural respect concerns whether an interaction is designed to succeed, not whether its outcome is morally acceptable or beneficial. It asks whether the operator has a stake in the user's successful navigation of the process, rather than whether the system is usable in the abstract. The following sections examine this mechanism in detail and explain why platforms occupy the position they do, and consider what follows for the legacy of UANTE.

Our concept of procedural respect aligns with Lon Fuller's account of the "internal morality of law" [13], which holds that the exercise of power becomes defective when its procedures are opaque, unintelligible, or insulated from revision. While platform recovery processes are not law, unilateral restrictions on access enacted through formalised procedures nonetheless function as a mode of governance. In this context, procedural failure does more than frustrate users: it subjects them to decisions they cannot understand, contest, or meaningfully navigate. This produces a structural paradox. Our ransomware operator must organise interaction through intelligible rules in order to survive, while the platform provider faces no comparable penalty for governing through a lack of clarity.

2 Why Platforms Cannot Be Clear

Our scenarios established our position; now, let us speculate on *what* permits this mechanism to operate in one case and not the other.

Based on our two scenarios, one significant pre-condition for the ransomware operator's design clarity becomes apparent. Platforms must sustain an appearance of acting in the user's interest a duty of care made visible through performance. Security messaging frames restrictions as protective interventions [27]; privacy policies claim safeguarding [11]; customer-service interfaces suggest responsiveness. These gestures maintain the fiction that the system acting *for* you by locking you out, flagging your behaviour, and otherwise interrupting your activity. In positioning restriction as protection the platform must appear as protector, and when something goes wrong, it is you who must be rendered blameworthy. When framed as such, a platform cannot admit that its fraud-reduction algorithm produced a false positive, or that resolving your case costs more than losing you as a user, or that your "suspicious activity" was indistinguishable from your ordinary behaviour. In such cases, being clear about what is happening would expose the gap between stated and actual purpose. Opacity preserves the fiction.

Further, we argue that this opacity fullfills a second, related function, and that is one of legitimation. By withholding reasons and timelines, platforms transmute contestable decisions into neutral system outcomes to foreclose scrutiny. This mirrors Citron's 'technological due process' failure [7], where systems retain the form of procedure while defeating its substance. Crucially, this opacity

exceeds what security necessity requires. While specific fraud-detection triggers may legitimately remain secret, the systematic absence of basic procedural information suggests institutional insulation rather than protective restraint. Procedural opacity is therefore not an accident of scale but a rational strategy for managing exposure where recovery success carries no operational value.

The ransomware operator faces no such constraint. Without pretence that the interaction serves the victim's interests, there remains no legitimation fiction to maintain. The power relation is explicit and transactional; for the victim, the interaction is simply unlucky to have occurred at all (and even this is sometimes acknowledged by the ransomware operator[4]). In terms of the perceived relationship, the ransomware victim retains agency in a meaningful sense: they can understand their situation, evaluate options (pay, refuse, seek technical remediation), and act. Of course such agency is coercively imposed, but within the bounds of our situation, victims are addressed as capable of comprehension and action.

We believe that the persistence of such a fiction, one that can survive market exposure, inherently must reveal the structural conditions that protect it, and we identify three such conditions. They are closely related, each blocking a different mechanism that would otherwise force platforms to design recovery processes that work. It represents the minimal set that jointly explains *why* procedural neglect is not an accident. The three conditions are: the platform's value extraction must not depend on recovery success, for otherwise operational necessity would have compelled a reformulation of this fictional relationship, users must lack a viable exit method, for otherwise competitive pressure would penalise the above behaviour, and comprehension must threaten rather than serve institutional interests, for otherwise transparency would be the cheaper strategy, as it has been within our ransomware example. The following section examines how each condition operates and why, together, they make such fictional relationships not merely possible but rational.

3 The Economics of Opacity

Platform revenue is frequently derived from user data, attention history, behavioural patterns, and network position, value generated through retention and monetisation, and not through any particular interaction succeeding. At the moment of account lockout, the platform has already extracted the value it seeks. Recovery therefore operates as a pure cost centre [2, 6]: every resource invested in helping users understand their situation, navigate resolution pathways, or receive meaningful feedback is an expense with no corresponding return. We term this condition *decoupled value*, the separation of platform revenue from interaction success. This creates organisational indifference to recovery quality that propagates to individuals who might otherwise profess commitment to the UANTE philosophy [9, 24]. When the organisation has no stake in recovery interaction quality, developers have no mandate (and no resources) to invest in it.

[4] As in the case of *Akira*, the ransomware example shown in Fig.1.

The ransomware operator, by contrast, cannot extract value until the interaction succeeds; payment requires navigation. Value and interaction are coupled, and that coupling produces investment. The same platform that guides new users through seamless onboarding abandons them to opacity during recovery. This internal contrast is explicable only by the economics of each interaction: onboarding converts prospects to revenue; recovery converts problems to costs. Identical organisations, identical users, radically different design investment.

Indifference alone, however, might be overcome by competitive pressure. Users poorly served by one platform could migrate to another. Yet network effects, data dependencies, and accumulated history create lock-in that makes switching prohibitively costly [19,21,25]. Even users who experience serious harm from opaque recovery processes are unlikely to abandon platforms where their social connections, professional identity, or transaction history reside. We term this condition *exit foreclosure*, [5] the structural prevention of competitive correction. This insulation from competitive consequence makes investment in recovery quality doubly unattractive, for not only does it generate no direct return, but failing to invest carries no penalty.

Freed from both internal incentive and external pressure, platforms face a third dynamic. User comprehension during recovery actively threatens institutional interests. A user who understands why they were locked out might contest the decision. A user who understands the appeals process might discover it consists of automated rejections. A user who understands the terms of service might recognise that the platform has disclaimed all responsibility for their situation [10]. Comprehension invites scrutiny; opacity forecloses it. We call this condition *opacity-as-protection*, which allows for the institutional benefit derived from user incomprehension.

We term this condition opacity-as-protection, the institutional benefit derived from systematic user incomprehension. While this opacity serves the legitimating function discussed in the previous section, its primary utility within this economic model is cost avoidance. Transparency creates actionable surfaces, specifically, specific reasons to contest, timelines to track, and decisions to appeal. Each of these generates operational friction. By withholding these details, opacity functions as a rigorous cost-control mechanism, preempting dispute resolution by rendering the user's situation navigable only through resignation.

To design for comprehension is, effectively, to design for vulnerability. Disclosing the specific trigger does not merely inform the user; it establishes grounds for substantive disagreement. Revealing the evidentiary basis might expose that the decision relied on thin, probabilistic signals rather than verified facts. Similarly, transparency regarding timelines would likely quantify, and thus expose, the total absence of service-level commitments. In this light, opacity prevents the user from seeing that the "appeal" is often a circular interaction with the same automated logic that issued the rejection. Silence protects the system not from

[5] Related accounts of platform power frame such conditions as the emergence of private procedural governance, where exit is formally available but substantively foreclosed; we do not develop these implications here.

attackers, but from the revelation of its own arbitrariness. Once arbitrariness is understood as something to be concealed rather than corrected, poor interaction quality in constrained-access scenarios appears not as negligence, but as an functional strategy. In combination, these conditions are self-reinforcing: sustained opacity normalises itself, eroding the user expectations that might otherwise generate pressure for change. Usable security research has typically framed opaque recovery interfaces as failures to apply known principles, yet the ransomware comparison allows us to suggest an alternative interpretation: organisations may understand the principles perfectly well and decline to implement them because user comprehension does not serve institutional interests.

Let us briefly consider what designing for comprehension would require. A platform enabling users to understand account recovery would need to disclose what triggered the restriction, what evidence supports the decision, what the user can do to resolve it, how long the process will take, and who is accountable if the process fails. Each disclosure creates exposure. The trigger might be contestable. The evidence might be thin or algorithmically generated without human review. The resolution pathway might reveal that appeals are processed by the same automated system that issued the original decision. The timeline might expose that the platform has no service-level commitment. Accountability might reveal that terms of service disclaim all responsibility.

Opacity protects against all of these exposures. Automated messages, generic denials, circular verification flows, and inaccessible escalation routes are not failures of design capacity but achievements of institutional self-protection. Stutzman and Hartzog proposed obscurity as a deliberate design choice to protect user privacy [35]; platforms have weaponised the same principle against users, deploying opacity to forestall accountability rather than to enable it [14,16,39][6].

4 Users Are Still Not the Enemy

Decoupled value, exit foreclosure, and opacity-as-protection: all of these conditions were nascent in 1999 but had not yet achieved their present dominance, as usable security increasingly falls into the hands of just a few large platform providers and third party services [20,34]. With such centralisation come deepened data dependencies, and the evolution of terms of service into instruments of comprehensive liability displacement. The philosophy of UANTE emerged in one structural environment and must now contend with another.

[6] Although not the focus of this paper, we must briefly mention here the parallel to dark patterns in consent interfaces, which is exceptional [5,22]. Dark patterns satisfy formal requirements, notices exist, boxes are clickable, yet systematically undermine substantive understanding. Cookie consent dialogs that require seventeen clicks to refuse but one click to accept [28, 37]; privacy policies written at reading levels that ensure incomprehension [36]; settings labyrinths that make opt-out technically possible but practically unreachable. These patterns share the same structural logic as opaque recovery processes: the form of user agency is preserved while its substance is defeated.

For those who adhere to that philosophy, our position yields a diagnostic question that cuts through professed values to a certain operational reality, where, when faced with usable security issues, we must ask: *does the operator need the interaction to succeed?* Where the answer is yes, expect investment in comprehensibility, actionability, tractability, and responsiveness—the dimensions through which procedural respect manifests. Where the answer is no, expect opacity, circularity, and indifference. The question is predictive precisely because it attends to structural position rather than declared intent. An organisation may profess user-centricity while its recovery interfaces betray the opposite; the diagnostic question explains why.

This framing also clarifies where intervention might prove effective. Conditions that create operational dependency on user success, outcome-based accountability, competitive pressure through portability, liability for interaction failures, would align institutional interest with user interest in ways that exhortation cannot. We do not develop these possibilities here, but the structural analysis points toward them.

Our position was established through a deliberately delimited comparison; ransomware and platform lockout represent cases where the structural conditions align starkly, and other user-security interactions may exhibit different configurations. Despite this, we believe our simple diagnostic question can generalise as a tool for identifying where procedural investment will and will not emerge.

Twenty-five years on, Sasse's body of work functions today less as a description of how systems behave and more as a commitment to how they *should*. In this paper, it provided the diagnostic vocabulary through which we can identify *where* that philosophy gains traction and *why* it fails to do so elsewhere. The structural conditions we have documented do not render the philosophy obsolete, rather they specify its conditions of application. Where operational dependency on user success can be created, through regulation, competitive pressure, or liability, the UANTE philosophy finds purchase. Where those conditions are absent, we should expect the opacity described in this paper. It should not be seen as fatalism, but as strategic clarity, for interventions that alter structural incentives will succeed where exhortation has failed. The users are still not the enemy; the task now is to make that truth consequential.

References

1. Adams, A., Sasse, M.A.: Users are not the enemy. Communications of the ACM **42**(12), 40–46 (1999)
2. Altinkemer, K., Wang, T.: Cost and benefit analysis of authentication systems. Decision Support Systems **51**(3), 394–404 (2011)
3. Beautement, A., Sasse, M.A., Wonham, M.: The compliance budget: Managing security behaviour in organisations. In: New Security Paradigms Workshop (NSPW). pp. 47–58 (2008)
4. Bell, G., Blythe, M., Sengers, P.: Making by making strange: Defamiliarization and the design of domestic technologies. ACM Transactions on Computer-Human Interaction **12**(2), 149–173 (2005)

5. Bösch, C., Erb, B., Kargl, F., Kopp, H., Pfattheicher, S.: Tales from the dark side: Privacy dark strategies and privacy dark patterns. Proceedings on Privacy Enhancing Technologies **2016**(4), 237–254 (2016)
6. Brostoff, S., Sasse, M.A.: 'ten strikes and you're out': Increasing the number of login attempts can improve password usability. In: Proceedings of the CHI 2003 Workshop on HCI and Security Systems (2003)
7. Citron, D.K.: Technological due process. Wash. UL Rev. **85**, 1249 (2007)
8. Connolly, L.Y., Wall, D.S.: The rise of crypto-ransomware in a changing cybercrime landscape: Taxonomising countermeasures. Computers & Security **87**, 101568 (2019)
9. Das Chowdhury, P., Hallett, J., Patnaik, N., Tahaei, M., Rashid, A.: Developers are neither enemies nor users: They are collaborators. In: 2021 IEEE Secure Development Conference (SecDev). pp. 47–55. IEEE (2021)
10. Das Chowdhury, P., Renaud, K., Rashid, A.: When data breaches happen, where does the buck stop? In: New Security Paradigms Workshop (NSPW). pp. 1–20 (2024)
11. Fernback, J., Papacharissi, Z.: Online privacy as legal safeguard: the relationship among consumer, online portal, and privacy policies. New Media & Society **9**(5), 715–734 (2007)
12. Filiz, B., Arief, B., Cetin, O., Hernandez-Castro, J.: On the effectiveness of ransomware decryption tools. Computers & Security **111**, 102469 (2021)
13. Fuller, L.L.: The Morality of Law. Yale University Press, New Haven (1969)
14. Gavazzi, A., Williams, R., Kirda, E., Lu, L., King, A., Davis, A., Leek, T.: A study of multi-factor and risk-based authentication availability. In: USENIX Security 2023. USENIX Association (2023)
15. George, B., Klems, M., Valeva, A.: A method for incorporating usable security into computer security courses. In: Proceedings of the 44th ACM Technical Symposium on Computer Science Education. pp. 681–686. ACM (2013)
16. Gray, C.M., Kou, Y., Battles, B., Hoggatt, J., Toombs, A.L.: The dark (patterns) side of ux design. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. pp. 1–14. ACM (2018)
17. Hadlington, L.J.: Exploring the psychological mechanisms used in ransomware splash screens. Tech. rep., De Montfort University (2017)
18. Hull, G., John, H., Arief, B.: Ransomware deployment methods and analysis: Views from a predictive model and human responses. Crime Science **8**, 1–22 (2019)
19. Krämer, J., Schnurr, D., Wohlfarth, M.: Winners, losers, and facebook: The role of social logins in the online advertising ecosystem. Management Science **65**(4), 1678–1699 (2019)
20. Kurtz, C., Burmeister, F.: Multi-role actors and rebounding effects across user interfaces-exploring big tech's privacy scandals and gdpr limitations in data ecosystems. In: International Conference on Human-Computer Interaction. pp. 283–303. Springer (2024)
21. Martin, K.: Platforms, privacy, and the honeypot problem. Harvard Journal of Law & Technology **37**, 1087 (2023)
22. Mathur, A., Mayer, J., Kshirsagar, M.: What makes a dark pattern... dark? design attributes, normative considerations, and measurement methods. In: CHI Conference on Human Factors in Computing Systems. ACM (2021)
23. Meland, P.H., Bayoumy, Y.F.F., Sindre, G.: The ransomware-as-a-service economy within the darknet. Computers & Security **92**, 101762 (2020)

24. Menges, U., Hielscher, J., Buckmann, A., Kluge, A., Sasse, M.A., Verret, I.: Why it security needs therapy. In: European Symposium on Research in Computer Security. pp. 335–356. Springer (2021)
25. Mody, M.A., Lu, L., Hanks, L.: 'it's not worth the effort!' examining service recovery in airbnb and other homesharing platforms. International Journal of Contemporary Hospitality Management **32**(9), 2991–3014 (2020)
26. Oz, H., Aris, A., Levi, A., Uluagac, A.S.: A survey on ransomware: Evolution, taxonomy, and defense solutions. ACM Computing Surveys **54**(11s), 1–37 (2022)
27. Rodríguez-Priego, N., Van Bavel, R., Vila, J., Briggs, P.: Framing effects on online security behavior. Frontiers in Psychology **11**, 527886 (2020)
28. Sanchez-Rola, I., Dell'Amico, M., Kotzias, P., Balzarotti, D., Bilge, L., Vervier, P.A., Santos, I.: Can i opt out yet? gdpr and the global illusion of cookie control. In: ACM AsiaCCS (2019)
29. Sasse, M.A.: Usability and trust in information systems. In: The Economics of Information Security. Edward Elgar (2005)
30. Sasse, M.A., Flechais, I.: Usable security: Why do we need it? how do we get it? In: Cranor, L.F., Garfinkel, S. (eds.) Security and Usability: Designing Secure Systems That People Can Use, pp. 13–30. O'Reilly (2005)
31. Sasse, M.A., Hielscher, J., Friedauer, J., Buckmann, A.: Rebooting it security awareness—how organisations can encourage and sustain secure behaviours. In: European Symposium on Research in Computer Security. pp. 248–265. Springer (2022)
32. Sharevski, F., Trowbridge, A., Westbrook, J.: Novel approach for cybersecurity workforce development: A course in secure design. In: 2018 IEEE Integrated STEM Education Conference (ISEC). pp. 175–180. IEEE (2018)
33. Shklovsky, V.: Art as technique. In: Lemon, L.T., Reis, M.J. (eds.) Russian Formalist Criticism: Four Essays, pp. 3–24. University of Nebraska Press (1965), originally published 1917
34. Stock, B., Johns, M., Steffens, M., Backes, M.: How the web tangled itself: Uncovering the history of {Client-Side} web ({In) Security}. In: 26th USENIX Security Symposium (USENIX Security 17). pp. 971–987 (2017)
35. Stutzman, F., Hartzog, W.: Obscurity by design. Washington Law Review **88**, 385–395 (2013)
36. Turow, J., Hennessy, M., Draper, N.: Persistent misperceptions: Americans' misplaced confidence in privacy policies, 2003–2015. Journal of Broadcasting & Electronic Media **62**(3), 461–478 (2018)
37. Utz, C., Degeling, M., Fahl, S., Schaub, F., Holz, T.: (un)informed consent: Studying gdpr consent notices in the field. In: ACM Conference on Computer and Communications Security (CCS). pp. 973–990 (2019)
38. Vidyaraman, S., Chandrasekaran, M., Upadhyaya, S.: Position: The user is the enemy. In: Proceedings of the 2007 Workshop on New Security Paradigms. pp. 75–80 (2008)
39. Wiefling, S., Lo Iacono, L., Dürmuth, M.: Is this really you? an empirical study on risk-based authentication applied in the wild. In: IFIP SEC 2019. pp. 134–148. Springer (2019)
40. Yilmaz, Y., Cetin, O., Arief, B., Hernandez-Castro, J.: Investigating the impact of ransomware splash screens. Journal of Information Security and Applications **61**, 102934 (2021)